arXiv:1703.05794v1 [stat.ME] 16 Mar 2017

Incorporating Covariates into Integrated Factor Analysis of Multi-View Data

Gen Li[1] and Sungkyu Jung[2]

[1]Department of Biostatistics, Mailman School of Public Health, Columbia University

[2]Department of Statistics, University of Pittsburgh

Abstract

In modern biomedical research, it is ubiquitous to have multiple data sets measured on the same set of samples from different views (i.e., multi-view data). For example, in genetic studies, multiple genomic data sets at different molecular levels or from different cell types are measured for a common set of individuals to investigate genetic regulation. Integration and reduction of multi-view data have the potential to leverage information in different data sets, and to reduce the magnitude and complexity of data for further statistical analysis and interpretation. In this paper, we develop a novel statistical model, called supervised integrated factor analysis (SIFA), for integrative dimension reduction of multi-view data while incorporating auxiliary covariates. The model decomposes data into joint and individual factors, capturing the joint variation across multiple data sets and the individual variation specific to each set respectively. Moreover, both joint and individual factors are partially informed by auxiliary covariates via nonparametric models. We devise a computationally efficient Expectation-Maximization (EM) algorithm to fit the model under some identifiability conditions. We apply the method to the Genotype-Tissue Expression (GTEx) data, and provide new insights into the variation decomposition of

gene expression in multiple tissues. Extensive simulation studies and an additional application to a pediatric growth study demonstrate the advantage of the proposed method over competing methods.

1 Introduction

In contemporary biomedical studies, researchers usually have access to multiple data sets for the same set of subjects from different views or heterogeneous sources. Such data are commonly referred to as multi-view data or multi-source data. For example, the Genotype-Tissue Expression (GTEx) project collects gene expression data from multiple human tissues for a common set of genotyped individuals to study genetic regulation (Ardlie et al., 2015). Different data sets may contain distinct but related information. It is important to understand the relations between variables in different sets, and leverage information across views for further statistical analysis such as inference, prediction and clustering. The process is often called data integration or data fusion.

Factor analysis is a popular tool for modeling dependence among multiple observed variables. It identifies a few latent factors that capture the majority of variation in data. The unknown factors and loadings in factor analysis are sometimes estimated via the principal component analysis (PCA). The obtained factors reduce the dimensionality of the original data and facilitate various statistical analyses. However, the conventional factor analysis only applies to a single data set. There is a pressing need for statistical methods that simultaneously identify the joint and individual structure in multiple data sets.

In addition to multiple primary data sets, auxiliary covariates are often collected on the same samples. In our motivating GTEx example, other than the gene expression data in multiple tissues, genotype data and experimental factors (e.g., batch effect) are also collected. These auxiliary data can be viewed as covariates, driving the underlying expression patterns in multiple tissues. Covariates are potential driving factors of the joint and individual structures in multi-view data. In other words, covariates provide *supervision* to the underlying patterns. Using covariates to inform the integration of multi-view data not only leads to accurate estimation of the

underlying patterns but also provides highly interpretable results.

In this paper, we develop a novel statistical framework called *Supervised Integrated Factor Analysis* (SIFA), for the integration and reduction of multi-view data informed by auxiliary covariates. SIFA decomposes multi-view data into low-rank joint structure and individual structure. It exploits a small number of joint factors to capture the shared patterns across all data sets, and separate individual factors to capture the specific patterns in each data set. Corresponding loading vectors identify the contribution of the variables to different factors. To allow auxiliary covariates to inform the latent structure, the model assumes each factor is potentially driven by the covariates and some random effects. We particularly consider regression models that flexibly accommodate parametric or nonparametric relations between factors and covariates. Through the regression models, the covariates exert supervision on the latent structure. We also extend the model to incorporate variable selection, in order to identify important covariates that drive different factors. Overall, SIFA provides a general framework for the covariate-driven factor analysis of multi-view data.

There is an extensive body of literature on the integrative analysis of multi-view data (Tseng et al., 2015). Here we particularly focus on data integration and reduction. Multiple factor analysis is an extension of the conventional factor analysis to multiple data sets (Abdi et al., 2013). The idea is to merge multiple data with weights and perform the factor analysis on the combined data. However, the method does not distinguish joint and individual structure and may lead to misleading results. More recently, new methods have been developed to decompose the total variation of multiple data sets into shared and individual variation (Löfstedt et al., 2013; Ray et al., 2014; Schouteden et al., 2014; Yang and Michailidis, 2016; Zhou et al., 2016). For example, Lock et al. (2013) adopts an iterative PCA approach to estimate the Joint and Individual Variation Explained (JIVE). However, a drawback of these methods is that they cannot take into account any auxiliary covariates in dimension reduction. When covariates are strongly associated with the latent structure of the multi-view data, incorporating the supervision effects from covariates promises to improve estimation accuracy and interpretability.

Recently, a couple of methods were proposed to allow covariates to inform factor analysis. Li et al. (2016) developed the Supervised Singular Value Decomposition

(SupSVD) method that exploits linear models to accommodate covariates in dimension reduction of a primary data matrix. Later, Fan et al. (2016) proposed the projected PCA that generalizes SupSVD by allowing nonparametric relations between covariates and factors. However, these methods are only suitable for a single data set, and cannot easily extend to multi-view data. To our best knowledge, there is no covariate-driven factor analysis method for multi-view data decomposition. Our proposed method will bridge the gap and provide a unified framework.

The rest of the paper is organized as follows. In Section 2, we propose a semiparametric latent variable model for SIFA and develop an Expectation-Maximization (EM) algorithm to fit the model. In Section 3, extensive simulation studies are conducted to compare the proposed method with existing methods. In Section 4, we apply SIFA to the GTEx multi-tissue genetic data to offer novel insights into the decomposition of genetic variation in a gene set across multiple tissues. In Section 5, we discuss possible directions for future research. Technical details, additional simulation results, and an application to the decoupled growth amplitude and phase data from the Berkeley Growth Study can be found in the online supplementary material.

2 Integrated Factor Analysis Framework

In this section, we first introduce the latent variable model for SIFA and discuss its connection to existing methods. Then we elaborate two sets of identifiability conditions, and devise model fitting algorithms under respective conditions. Finally we propose rank selection methods to determine the joint and individual ranks in the model.

2.1 Model

Let $\boldsymbol{Y}_1, \cdots, \boldsymbol{Y}_K$ be K primary data matrices of size $n \times p_1, \cdots, n \times p_K$ for the same set of samples collected from K different sources. Each row corresponds to a sample and each column is a variable. Let $\boldsymbol{X}$ be an $n \times q$ data matrix containing covariates for the matched samples. The goal is to identify low-rank joint and individual patterns from the primary data matrices while accounting for the supervision effects from the covariates. Without loss of generality, we center each column of the primary data

and the covariates to remove the mean effect of each variable.

We propose a latent variable model called SIFA for the integrative factor analysis of multiple data matrices. For $k = 1, \cdots, K$, the SIFA model is as follows (without special notice, the index k takes integer values from 1 to K):

$$\begin{aligned} \boldsymbol{Y}_k &= \boldsymbol{J}_k + \boldsymbol{A}_k + \boldsymbol{E}_k, && (1) \\ \boldsymbol{J}_k &= \boldsymbol{U}_0 \boldsymbol{V}_{0,k}^T, && (2) \\ \boldsymbol{A}_k &= \boldsymbol{U}_k \boldsymbol{V}_k^T, && (3) \\ \boldsymbol{U}_0 &= \boldsymbol{f}_0(\boldsymbol{X}) + \boldsymbol{F}_0, && (4) \\ \boldsymbol{U}_k &= \boldsymbol{f}_k(\boldsymbol{X}) + \boldsymbol{F}_k. && (5) \end{aligned}$$

In (1), we adopt a signal-plus-noise model to capture the important patterns in each data set. This type of model is commonly used in the dimension reduction literature (cf. Shabalin and Nobel, 2013). More specifically, the data matrix $\boldsymbol{Y}_k$ consists of signal $\boldsymbol{J}_k + \boldsymbol{A}_k$ and noise $\boldsymbol{E}_k$. The matrix $\boldsymbol{J}_k$ captures the joint structure shared across multiple sources, and the matrix $\boldsymbol{A}_k$ captures the individual structure specific to this data source. The noise matrix $\boldsymbol{E}_k$ is assumed to have independent and identically distributed (i.i.d.) entries from a normal distribution $\mathcal{N}(0, \sigma_k^2)$, where σ_k^2 is an unknown parameter.

In (2) and (3), we assume that the joint and individual patterns for $\boldsymbol{Y}_k$ have low-rank decomposition. Let r_0 be the underlying rank of the joint structure and r_k be the rank of the individual structure in the kth data set. Correspondingly, $\boldsymbol{U}_0$ and $\boldsymbol{U}_k$ are $n \times r_0$ and $n \times r_k$ (latent) factor matrices, and $\boldsymbol{V}_{0,k}$ and $\boldsymbol{V}_k$ are $p_k \times r_0$ and $p_k \times r_k$ loading matrices. In particular, $\boldsymbol{U}_0$ contains r_0 joint factors shared across different data sets, and $\boldsymbol{V}_0 = (\boldsymbol{V}_{0,1}^T, \cdots, \boldsymbol{V}_{0,K}^T)^T$ contains r_0 corresponding joint loadings. The matrices $\boldsymbol{U}_k$ and $\boldsymbol{V}_k$ contain r_k individual factors and loadings respectively. Following the convention of the factor analysis, we assume the factors are independent and the loadings are orthonormal within each matrix. Namely, $\boldsymbol{V}_k^T \boldsymbol{V}_k = \boldsymbol{I}_{r_k}$ for each $k = 0, 1, \cdots, K$, where $\boldsymbol{I}_{r_k}$ denotes the $r_k \times r_k$ identity matrix (we shall drop the subscript when it does not cause any confusion).

In order to capture the driving effects of covariates on the low-rank structure, we propose to regress the latent factors on the covariates via (4) and (S.12). The mapping functions $\boldsymbol{f}_k(\cdot) : \mathbb{R}^q \mapsto \mathbb{R}^{r_k}$ $(k = 0, 1, \cdots, K)$ are unknown parametric

or nonparametric functions. With a slight abuse of notation, we use $\boldsymbol{f}_k(\boldsymbol{X})$ $(k = 0, 1, \cdots, K)$ to represent row-wise mappings. Namely, $\boldsymbol{f}_k(\boldsymbol{X})$ is an $n \times r_k$ matrix whose ith row corresponds to $\boldsymbol{f}_k(\boldsymbol{x}_{(i)})$, where $\boldsymbol{x}_{(i)}$ is the ith row of $\boldsymbol{X}$ $(i = 1, \cdots, n)$. The mapping functions capture flexible relations between covariates and the latent factors. In practice, users can determine whether to use nonparametric functions or parametric functions (e.g., linear functions). Any unknown variation in the factors is contained in the random matrices $\boldsymbol{F}_k$ $(k = 0, 1, \cdots, K)$. In particular, we assume each row of $\boldsymbol{F}_k$ follows an i.i.d. multivariate normal distribution with zero mean and an unknown diagonal covariance matrix $\boldsymbol{\Sigma}_k$ (with positive, distinct, and decreasing diagonal values). Moreover, we assume $\boldsymbol{F}_0$, $\boldsymbol{F}_k$'s, and $\boldsymbol{E}_k$'s are mutually independent.

The proposed SIFA model provides a general framework for the factor analysis of multi-view data. After rearranging the formulas, we get an equivalent form of the model as

$$\boldsymbol{Y}_k = \boldsymbol{f}_0(\boldsymbol{X})\boldsymbol{V}_{0,k}^T + \boldsymbol{f}_k(\boldsymbol{X})\boldsymbol{V}_k^T + \boldsymbol{F}_0\boldsymbol{V}_{0,k}^T + \boldsymbol{F}_k\boldsymbol{V}_k^T + \boldsymbol{E}_k. \tag{6}$$

It is easy to see that the SIFA model decomposes the kth data matrix $\boldsymbol{Y}_k$ into five parts: 1) $\boldsymbol{f}_0(\boldsymbol{X})\boldsymbol{V}_{0,k}^T$ is the joint deterministic structure (because $\boldsymbol{f}_0(\boldsymbol{X})$ is shared across multiple data sources and non-random) driven by the covariates; 2) $\boldsymbol{f}_k(\boldsymbol{X})\boldsymbol{V}_k^T$ is the individual deterministic structure; 3) $\boldsymbol{F}_0\boldsymbol{V}_{0,k}^T$ is the joint random structure capturing any unknown variation; 4) $\boldsymbol{F}_k\boldsymbol{V}_k^T$ is the individual random structure; 5) $\boldsymbol{E}_k$ is the random noise. With proper identifiability conditions which we will discuss later, the SIFA model attributes the total variation to different parts. Different model components will facilitate different analyses. For example, the joint factors in $\boldsymbol{f}_0(\boldsymbol{X}) + \boldsymbol{F}_0$ can be potentially used for consensus clustering; the individual loadings in $\boldsymbol{V}_k$ can be used to investigate the dependence among variables in the kth data source.

We remark that the proposed SIFA model (6) subsumes many existing methods as special cases. When $K = 1$, i.e., with only one primary data set $\boldsymbol{Y}$, there is no distinction between the joint structure and the individual structure. Consequently, the model degenerates to

$$\boldsymbol{Y} = (\boldsymbol{f}(\boldsymbol{X}) + \boldsymbol{F})\boldsymbol{V}^T + \boldsymbol{E},$$

which corresponds to the projected PCA model proposed by Fan et al. (2016). In

particular, if we let the function $\boldsymbol{f}(\cdot)$ be a linear mapping, i.e., $\boldsymbol{f}(\boldsymbol{X}) = \boldsymbol{X}\boldsymbol{B}$, where $\boldsymbol{B}$ is a $q \times r$ coefficient matrix, the above model further connects to the SupSVD model developed in Li et al. (2016). Furthermore, if we eliminate the covariate supervision by setting $\boldsymbol{f}(\boldsymbol{X}) = \boldsymbol{0}$, the model degenerates to the conventional factor analysis model or the probabilistic PCA model (Tipping and Bishop, 1999). When $K \geq 2$, without accounting for the covariates (i.e., $\boldsymbol{f}_k(\boldsymbol{X}) = \boldsymbol{0}; k = 0, 1, \cdots, K$), the SIFA model reduces to

$$\boldsymbol{Y}_k = \boldsymbol{F}_0 \boldsymbol{V}_{0,k}^T + \boldsymbol{F}_k \boldsymbol{V}_k^T + \boldsymbol{E}_k.$$

This coincides with the JIVE model (Lock et al., 2013) if we assume $\boldsymbol{F}_0$ and $\boldsymbol{F}_k$ are unknown score matrices for the joint and individual structures. The SIFA model unifies and generalizes the above models, and provides a general framework for the integration and reduction of multi-view data informed by covariates.

2.2 Identifiability

Suppose $\theta_0 = \{\boldsymbol{f}_0(\cdot), \boldsymbol{f}_k(\cdot), \boldsymbol{V}_0, \boldsymbol{V}_k, \boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_k, \sigma_k^2; k = 1, \ldots, K\}$ is a parameter set for Model (6), satisfying the basic conditions of $\boldsymbol{V}_k^T \boldsymbol{V}_k = \boldsymbol{I}$ and $\boldsymbol{\Sigma}_k$ being diagonal with distinct (decreasing) positive diagonal values for each $k = 0, 1, \cdots, K$. If there is only one primary data set (i.e., $K = 1$), the model is uniquely defined (Li et al., 2016). However, when there are multiple data sets, the above basic conditions are no longer sufficient for identifiability.

To be specific, let Θ be the collection of parameter sets θ satisfying the basic conditions and having equal likelihood $\mathcal{L}(\boldsymbol{Y}_1, \cdots, \boldsymbol{Y}_K \mid \theta)$ (defined later in (7)) with θ_0 for any data $\boldsymbol{Y}_1, \cdots, \boldsymbol{Y}_K$. Namely, any parameter set $\theta \in \Theta$ and θ_0 are *observationally equivalent* for Model (6), i.e., Θ is the *equivalence class* of θ_0. We note that the collection Θ typically contains multiple elements (see the supplementary material for examples of some equivalent models). In other words, θ_0 is unidentifiable. This non-identifiability is mainly caused by the indistinguishable individual and joint structures. Different elements in Θ may have different sets of ranks, or the same set of ranks but different parameters. Additional regularity conditions are needed to enforce the identifiability of the SIFA model. For this, we propose two sets of sufficient conditions.

First, we consider a set of *general conditions* for each $k = 1, \cdots, K$:

A1. Each submatrix $\boldsymbol{V}_{0,k}$ of the joint loading matrix $\boldsymbol{V}_0$ has full column rank;

A2. The columns in $\boldsymbol{V}_{0,k}$ and $\boldsymbol{V}_k$ are linearly independent, and $r_0 + r_k < p_k$.

Loosely speaking, Condition A1 guarantees that the joint loading matrix $\boldsymbol{V}_0$ indeed captures the joint structure, and does not contain any structure only pertaining to a subset of the K data sets. Condition A2 ensures that the joint and individual patterns are well separated, and are not interchangeable. With both conditions, Model (6) is identifiable as shown in the following proposition (the proof is postponed to the supplementary material).

Proposition 2.1. *Let θ_0 be a parameter set satisfying Conditions A1 and A2. For any element θ in the equivalent class Θ of θ_0, if θ also satisfies Conditions A1 and A2, then θ is equal to θ_0 up to trivial sign changes. Moreover, by writing $r_0(\theta)$ as the rank of $\boldsymbol{V}_0$ in the parameter set θ, we have $r_0(\theta_0) \leq r_0(\theta)$ for all $\theta \in \Theta$.*

The proposition guarantees that the general conditions are sufficient for model identifiability. The identifiability is defined up to trivial column-wise sign changes in $\boldsymbol{V}_k$ and $\boldsymbol{U}_k$ ($k = 0, 1, \cdots, K$). In practice, one could easily fix the signs by setting the first nonzero entry of each column of $\boldsymbol{V}_k$ to be positive. Correspondingly, the sign of each column of $\boldsymbol{U}_k$ is fixed.

Remark: Technically, the general conditions may slightly affect the generalizability of the model. Condition A1 rules out the possibility of any partially joint structure pertaining to multiple but not all data sets. Namely, the model cannot identify common patterns across a subset of data sets. The same issue exists for JIVE as well. This is a future research direction as discussed in Section 5. Nevertheless, in practice, the general conditions are suitable for most applications.

In some circumstances, it is desired to further restrict the model parameters for better interpretation and computation. We consider the following *orthogonal conditions*:

B1. The columns of $\boldsymbol{V}_{0,k}$ are orthogonal with norm $1/\sqrt{K}$, i.e., $\boldsymbol{V}_{0,k}^T \boldsymbol{V}_{0,k} = \frac{1}{K}\boldsymbol{I}$;

B2. The columns in $\boldsymbol{V}_{0,k}$ and $\boldsymbol{V}_k$ are orthogonal ($\boldsymbol{V}_{0,k}^T \boldsymbol{V}_k = \boldsymbol{0}$), and $r_0 + r_k < p_k$.

Apparently, Conditions B1 and B2 are sufficient conditions for Conditions A1 and A2. Therefore, they are also sufficient conditions for the identifiability of the SIFA model. Condition B1 implies that different data sets contribute roughly equally to the joint factors $\boldsymbol{U}_0$ (i.e., columns in $\boldsymbol{V}_{0,k_1}$ and $\boldsymbol{V}_{0,k_2}$ have the same ℓ_2 norm, for $k_1 \neq k_2$). In many real applications (e.g., the GTEx example in Section 4), when the data are properly preprocessed, the equal contribution assumption can be well justified. Conditions B1 and B2 also imply that the combined loadings $(\boldsymbol{V}_{0,k}, \boldsymbol{V}_k)$ for the kth data set are mutually orthogonal. For high dimensional data, it is reasonable to assume that the orthogonality between different loadings holds (Ahn and Marron, 2010). When both assumptions are justified, it is beneficial to study the SIFA model under the orthogonal conditions. These conditions not only improve model interpretation, but also facilitate computation as discussed in the next subsection.

Remark: The SIFA model with the general conditions is equivariant under individual scaling of each data set. In other words, at the population level, the model is not affected by weighing multiple data sets differently. In practice, to avoid numerical instability, it is recommended to normalize different data sets to the same scale before estimation (e.g., set the Frobenius norm of every data set to be 1). The orthogonal conditions do not have the equivariant property under rescaling. Thus, if the orthogonal assumptions are justifiable, one should directly apply the method without scaling the data. See the supplementary material for more details. There, we also discuss the effect of imbalanced dimensions of different data sets.

2.3 Algorithm

To estimate the model parameters in $\theta_0 = \{\boldsymbol{f}_0(\cdot), \boldsymbol{f}_k(\cdot), \boldsymbol{V}_{0,k}, \boldsymbol{V}_k, \boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_k, \sigma_k^2; k = 1, \ldots, K\}$ for Model (6), we use a maximum likelihood approach. We assume all random variables are from normal distributions. For the ease of presentation, $\boldsymbol{V}_\star = \text{blkdiag}(\boldsymbol{V}_1, \cdots, \boldsymbol{V}_K)$ denotes the combined individual loading matrix of size $\sum_{k=1}^K p_k \times \sum_{k=1}^K r_k$, which is a block-wise diagonal matrix with K diagonal blocks $\boldsymbol{V}_1, \cdots, \boldsymbol{V}_K$. We also let $\boldsymbol{U}_\star = (\boldsymbol{U}_1, \cdots, \boldsymbol{U}_K) = (\boldsymbol{f}_1(\boldsymbol{X}) + \boldsymbol{F}_1, \cdots, \boldsymbol{f}_K(\boldsymbol{X}) + \boldsymbol{F}_K)$ denote the combined individual factor matrix. Let $\boldsymbol{Y}_\star = (\boldsymbol{Y}_1, \cdots, \boldsymbol{Y}_K)$ and $\boldsymbol{E}_\star = (\boldsymbol{E}_1, \cdots, \boldsymbol{E}_K)$ be the combined primary data matrix and noise matrix respectively. As a result, the

SIFA model can be succinctly expressed as

$$\boldsymbol{Y}_\star = \boldsymbol{U}_0\boldsymbol{V}_0^T + \boldsymbol{U}_\star\boldsymbol{V}_\star^T + \boldsymbol{E}_\star.$$

The log likelihood function of the SIFA model is

$$\log \mathcal{L}(\boldsymbol{Y}_\star \mid \theta_0) = \sum_{i=1}^{n}\left[-\frac{\sum_{k=1}^{K} p_k}{2}\log 2\pi - \frac{1}{2}\log|\boldsymbol{\Sigma}_\star| - \frac{1}{2}(\boldsymbol{y}_{\star(i)} - \boldsymbol{\mu}_{\star(i)})^T\boldsymbol{\Sigma}_\star^{-1}(\boldsymbol{y}_{\star(i)} - \boldsymbol{\mu}_{\star(i)})\right], \quad (7)$$

where $\boldsymbol{y}_{\star(i)}$ is a column vector corresponding to the ith row of $\boldsymbol{Y}_\star$, and $\boldsymbol{\mu}_{\star(i)}$ and $\boldsymbol{\Sigma}_\star$ are the marginal expectation and covariance matrix of $\boldsymbol{y}_{\star(i)}$ respectively. In particular,

$$\boldsymbol{\mu}_{\star(i)}^T = \boldsymbol{f}_0(\boldsymbol{x}_{(i)})\boldsymbol{V}_0^T + \left[\boldsymbol{f}_1(\boldsymbol{x}_{(i)})\boldsymbol{V}_1^T, \cdots, \boldsymbol{f}_K(\boldsymbol{x}_{(i)})\boldsymbol{V}_K^T\right],$$

where $\boldsymbol{x}_{(i)}$ is a column vector corresponding to the ith row of $\boldsymbol{X}$, and $\boldsymbol{f}_k(\boldsymbol{x}_{(i)})$ is a row vector of length r_k $(k = 0, 1, \cdots, K)$. The grand covariance matrix $\boldsymbol{\Sigma}_\star$ has the form

$$\boldsymbol{\Sigma}_\star = \boldsymbol{V}_0\boldsymbol{\Sigma}_0\boldsymbol{V}_0^T + \boldsymbol{V}_\star\boldsymbol{\Sigma}_{\boldsymbol{F}}\boldsymbol{V}_\star^T + \boldsymbol{\Sigma}_{\boldsymbol{E}},$$

where $\boldsymbol{\Sigma}_{\boldsymbol{F}} = \text{blkdiag}(\boldsymbol{\Sigma}_1, \cdots, \boldsymbol{\Sigma}_K)$ and $\boldsymbol{\Sigma}_{\boldsymbol{E}} = \text{blkdiag}(\sigma_1^2\boldsymbol{I}_{r_1}, \cdots, \sigma_k^2\boldsymbol{I}_{r_K})$. The optimization of the above log likelihood function under the identifiability conditions is computationally prohibitive because the likelihood function involves unknown nonparametric functions and the conditions are non-convex.

To circumvent the computational issue, we resort to the hierarchical form of the SIFA model in (1)–(S.12) and treat $\boldsymbol{U}_0$ and $\boldsymbol{U}_\star$ as latent variables, and derive an *Expectation-Maximization* (EM) algorithm. Specifically, in the E step, we calculate the conditional distribution of the latent variables $(\boldsymbol{U}_0, \boldsymbol{U}_\star)$ given the data $\boldsymbol{Y}_\star$ and the previously estimated model parameters. In the M step, we maximize the conditional expectation of the joint log likelihood of the latent variables and the data. The joint log likelihood is partitioned into the log likelihood of $(\boldsymbol{U}_0, \boldsymbol{U}_\star)$ and the conditional log likelihood of $\boldsymbol{Y}_\star$ given $(\boldsymbol{U}_0, \boldsymbol{U}_\star)$. Furthermore, since the latent variables $\boldsymbol{U}_0, \boldsymbol{U}_1, \cdots, \boldsymbol{U}_K$ are mutually independent, the log likelihood of $(\boldsymbol{U}_0, \boldsymbol{U}_\star)$ is further partitioned. Consequently, the M step is to solve the following problems under the respective identifiability conditions:

$$\max_{\boldsymbol{f}_k(\cdot), \boldsymbol{\Sigma}_k} \quad \mathbb{E}_{\boldsymbol{U}_k|\boldsymbol{Y}_\star}\,\mathcal{L}(\boldsymbol{U}_k), \quad k = 0, 1, \cdots, K, \quad (8)$$

$$\max_{\boldsymbol{V}_0, \boldsymbol{V}_\star, \sigma_1^2, \cdots, \sigma_K^2} \quad \mathbb{E}_{\boldsymbol{U}_0, \boldsymbol{U}_\star|\boldsymbol{Y}_\star}\,\mathcal{L}(\boldsymbol{Y}_\star|\boldsymbol{U}_0, \boldsymbol{U}_\star), \quad (9)$$

where $\mathbb{E}_{\boldsymbol{U}_0,\boldsymbol{U}_\star|\boldsymbol{Y}_\star}(\cdot)$ represents the conditional expectation with respect to $(\boldsymbol{U}_0, \boldsymbol{U}_\star)$. For simplicity, hereafter we will use $\mathbb{E}(\cdot)$ to denote the conditional expectations. Below we shall outline the key steps of the M step. More details can be found in the supplementary material.

It can be shown that in (8) each entry of the vector-valued function $\boldsymbol{f}_k(\cdot) = (f_{k,1}(\cdot), \cdots, f_{k,r_k}(\cdot))$ can be separately estimated via solving a least square problem

$$\widehat{f_{k,j}(\cdot)} = \underset{f_{k,j}(\cdot)}{\arg\min} \ \|\mathbb{E}(\boldsymbol{u}_{k,j}) - f_{k,j}(\boldsymbol{X})\|_{\mathbb{F}}^2, \quad j = 1, \cdots, r_k; \ k = 0, 1, \cdots, K, \tag{10}$$

where $\boldsymbol{u}_{k,j}$ is the jth column of $\boldsymbol{U}_k$, and $\|\cdot\|_{\mathbb{F}}$ denotes the Frobenius norm. If $f_{k,j}(\cdot)$ is a parametric function, the above problem can be solved via a Newton-Raphson method. In particular, if linear, it is explicitly solved via the ordinary least squares. If $f_{k,j}(\cdot)$ is nonparametric, the problem becomes nonparametric regression. Standard kernel methods and spline-based methods can be readily applied here (cf. Fan and Gijbels, 1996; Hollander et al., 2013). When the dimension of the covariates is high, we can assume $f_{k,j}(\cdot)$ to be an additive model and easily incorporate variable selection through penalized methods (Tibshirani, 1996; Ravikumar et al., 2009). To sum up, regardless of the forms of the functions, $\{\boldsymbol{f}_k(\cdot)\}_{k=0,\cdots,K}$ can be easily estimated using existing methods.

Subsequently, it is easy to obtain the closed-form optimizer of (8) with respect to $\boldsymbol{\Sigma}_k$ as:

$$\widehat{\boldsymbol{\Sigma}_k} = \frac{1}{n}\mathrm{diag}\left\{\mathbb{E}\left[\left(\boldsymbol{U}_k - \widehat{\boldsymbol{f}_k}(\boldsymbol{X})\right)^T\left(\boldsymbol{U}_k - \widehat{\boldsymbol{f}_k}(\boldsymbol{X})\right)\right]\right\}, \ k = 0, 1, \cdots, K,$$

where $\mathrm{diag}(\boldsymbol{S})$ is the diagonal matrix consisting of the diagonal values of $\boldsymbol{S}$, and $\widehat{\boldsymbol{f}_k}(\cdot)$'s are the estimated covariate functions.

From (9), we obtain the estimates of the loading matrices and the noise variances under different identifiability conditions.

Under the general conditions A1 and A2, there are no explicit solutions of (9) for $\boldsymbol{V}_0$ and $\boldsymbol{V}_\star$. We propose to iteratively update the estimates of the loading matrices in a block-wise coordinate descent fashion. In particular, we cycle through the following steps: given $\boldsymbol{V}_0$, update $\boldsymbol{V}_k$'s in parallel via solving

$$\min_{\boldsymbol{V}_k:\boldsymbol{V}_k^T\boldsymbol{V}_k=\boldsymbol{I}} \quad \mathbb{E}\|\boldsymbol{Y}_k - \boldsymbol{U}_0\boldsymbol{V}_{0,k}^T - \boldsymbol{U}_k\boldsymbol{V}_k^T\|_{\mathbb{F}}^2; \tag{11}$$

and given $\boldsymbol{V}_\star$, update $\boldsymbol{V}_0$ by solving

$$\min_{\boldsymbol{V}_0:\boldsymbol{V}_0^T\boldsymbol{V}_0=\boldsymbol{I}} \quad \sum_{k=1}^{K} \sigma_k^{-2}\mathbb{E}\|\boldsymbol{Y}_k - \boldsymbol{U}_k\boldsymbol{V}_k^T - \boldsymbol{U}_0\boldsymbol{V}_{0,k}^T\|_{\mathbb{F}}^2. \tag{12}$$

It can be shown that the optimization problem (11) has a closed-form solution $\widehat{\boldsymbol{V}_k} = \boldsymbol{L}\boldsymbol{R}^T$, where $\boldsymbol{L}$ and $\boldsymbol{R}$ contain the left and right singular vectors of $\boldsymbol{Y}_k^T\mathbb{E}(\boldsymbol{U}_k) - \boldsymbol{V}_{0,k}\mathbb{E}(\boldsymbol{U}_0^T\boldsymbol{U}_k)$. The optimization (12) does not have an analytical solution due to the possibly different σ_k^2's. As a remedy, we relax the orthogonality constraint in (12) temporarily, and obtain an intermediate estimator of $\boldsymbol{V}_{0,k}$ as

$$\widetilde{\boldsymbol{V}_{0,k}} = \left[\boldsymbol{Y}_k^T\mathbb{E}(\boldsymbol{U}_0) - \boldsymbol{V}_k\mathbb{E}(\boldsymbol{U}_k^T\boldsymbol{U}_0)\right]\left[\mathbb{E}(\boldsymbol{U}_0^T\boldsymbol{U}_0)\right]^{-1}.$$

To impose the orthogonality constraint, the final estimator of $\boldsymbol{V}_0$ is the eigenvectors of $\widetilde{\boldsymbol{V}_0}\widehat{\boldsymbol{\Sigma}_0}\widetilde{\boldsymbol{V}_0}^T$. Correspondingly, we update the diagonal values of $\widehat{\boldsymbol{\Sigma}_0}$ to be the eigenvalues of $\widetilde{\boldsymbol{V}_0}\widehat{\boldsymbol{\Sigma}_0}\widetilde{\boldsymbol{V}_0}^T$. This additional standardization step ensures that $\boldsymbol{\Sigma}_\star$ in the likelihood function (7) remains unchanged. A similar approach was used in Li et al. (2016). As a result, the loading matrices are estimated under the general conditions. We remark that in practice, a one-step update in each EM iteration is usually accurate enough and there is no need to iterate.

Under the orthogonal conditions B1 and B2, the computation can be greatly simplified. The loading matrices $\boldsymbol{V}_{0,k}$ and $\boldsymbol{V}_k$ can be estimated simultaneously and explicitly. By writing $\boldsymbol{W}_k = (\sqrt{K}\boldsymbol{V}_{0,k}, \boldsymbol{V}_k)$, the optimization (9) is equivalent to

$$\min_{\boldsymbol{W}_k:\boldsymbol{W}_k^T\boldsymbol{W}_k=\boldsymbol{I}} \left\|\boldsymbol{Y}_k - \left(\frac{1}{\sqrt{K}}\mathbb{E}(\boldsymbol{U}_0), \mathbb{E}(\boldsymbol{U}_k)\right)\boldsymbol{W}_k^T\right\|_{\mathbb{F}}^2,$$

which is exactly an orthogonal Procrustes problem (Gower and Dijksterhuis, 2004). The optimizer has the explicit expression as $\widehat{\boldsymbol{W}_k} = (\sqrt{K}\widehat{\boldsymbol{V}_{0,k}}, \widehat{\boldsymbol{V}_k}) = \boldsymbol{L}\boldsymbol{R}^T$, where $\boldsymbol{L}$ and $\boldsymbol{R}$ contain the left and right singular vectors of $\boldsymbol{Y}_k^T\left(1/\sqrt{K}\mathbb{E}(\boldsymbol{U}_0), \mathbb{E}(\boldsymbol{U}_k)\right)$. Subsequently, it is easy to decouple $\widehat{\boldsymbol{V}_{0,k}}$ and $\widehat{\boldsymbol{V}_k}$, and obtain closed-form estimators for different loading matrices.

Once the loading matrices are estimated, solving (9) with respect to σ_k^2's, we obtain the closed-form optimizers as:

$$\widehat{\sigma_k^2} = \frac{1}{np_k}\mathbb{E}\|\boldsymbol{Y}_k - \boldsymbol{U}_k\widehat{\boldsymbol{V}_k}^T - \boldsymbol{U}_0\widehat{\boldsymbol{V}_{0,k}}^T\|_{\mathbb{F}}^2,\ k = 1, \cdots, K.$$

A step-by-step description of the algorithm can be found in the supplementary material.

2.4 Rank Selection

Up to now, we assume the ranks for the joint and individual structures in the SIFA model are known. In practice, we often need to estimate the ranks from data. The choice of the ranks is crucial for parameter estimation and model interpretation. In total, there are $K+1$ ranks to be determined. Here we propose a two-step procedure to get a crude estimate of the ranks, and an optional likelihood cross validation (LCV) method for refining the estimate.

Since Model (6) can be viewed as a special form of a signal-plus-noise model, a natural first step is to estimate the rank of the underlying signal of each data set $\boldsymbol{Y}_k$ (denoted as $r_k^\star$) and the rank of the underlying signal of the combined data set $\boldsymbol{Y}_\star$ (denoted as $r_{total}^\star$) respectively. There are many existing methods to this purpose, such as the scree plot, the total variance explained criterion, hypothesis testing methods. Users can choose their favorite methods. Once estimated, we use $r_k^\star$ and $r_{total}^\star$ to calculate the ranks for different structures in Model (6). Under either set of identifiability conditions, the following equations hold:

$$r_{total}^\star = r_0 + \sum_{k=1}^{K} r_k, \quad r_k^\star = r_0 + r_k,$$

for $k = 1, \cdots, K$, where $r_0, r_1, \cdots, r_K$ are the joint and individual ranks for the SIFA model. Solving the equation system, we get

$$r_0 = \frac{\sum_{k=1}^{K} r_k^\star - r_{total}^\star}{K-1}, \quad r_k = r_k^\star - r_0,$$

which serve as good initial estimators of the ranks. Numerically, the estimate of r_0 may be non-integer or even negative when $K > 2$. In that case, we suggest rounding the estimate to the nearest non-negative integer. Then we plug it into the second equation to get an estimate of r_k. If the estimate of r_k is negative, it can be set to 0.

The above two-step procedure typically provides a good initial estimate of the ranks. If it is desired to further refine the rank estimation, one may exploit a more computationally intensive N-fold LCV approach. The idea is to randomly split the samples into N groups across different data sets. In each run, we withhold one group as the testing set and use the remaining $N-1$ groups as the training set to fit Model (6) with different sets of ranks. For each set of ranks, the corresponding

LCV score is the value of negative log likelihood, evaluated using (7) on the testing data. We repeat the procedure N times, and choose the set of ranks corresponding to the smallest average LCV score. A more detailed description can be found in the supplementary material.

3 Simulation Studies

In this section, we conduct comprehensive simulation studies to demonstrate the advantage of the proposed methods. We compare SIFA (under both sets of identifiability conditions) with JIVE (the original version and a covariate-adjusted version, denoted by cov-JIVE), SupSVD, and PCA. For cov-JIVE, we first regress different data sets on the covariates, and then apply JIVE to the residuals.

3.1 Simulation Settings

We consider two primary data sets $\boldsymbol{Y}_1$ and $\boldsymbol{Y}_2$ (i.e., $K = 2$) on the same set of samples with sample size $n = 500$, and dimension $p_1 = p_2 = 200$. The data are simulated from Model (6) with different parameters. We first consider 3 settings where, loosely speaking, the generative models are JIVE, SIFA under the general conditions (denoted as SIFA-A), and SIFA under the orthogonal conditions (denoted as SIFA-B). In particular, the SIFA-A and SIFA-B models employ linear models between covariates (with dimension $q = 10$) and latent factors. The true ranks of the joint and individual patterns are $r_0 = 2$, $r_1 = r_2 = 3$. Some important features of these settings are described below.

- **Setting 1** (JIVE Model): For $k = 0, 1, 2$, the factors in $\boldsymbol{U}_k$ are randomly generated and mutually independent (with $\boldsymbol{f}_k(\cdot) = \boldsymbol{0}$ in (6)); the loadings in $\boldsymbol{V}_k$ and the covariance $\boldsymbol{\Sigma}_k$ satisfy the basic conditions. The measurement errors in $\boldsymbol{E}_1$ and $\boldsymbol{E}_2$ are i.i.d. with different variances. To test the robustness of the proposed method, we randomly generate 10 covariates unrelated with the factors, and incorporate them in the SIFA estimation.

- **Setting 2** (SIFA-A Model): The joint and individual factors are generated from the linear model $\boldsymbol{U}_k = \boldsymbol{X}\boldsymbol{B}_k + \boldsymbol{F}_k$ for $k = 0, 1, 2$. The loadings in $\boldsymbol{V}_0, \boldsymbol{V}_1$

and $\boldsymbol{V}_2$ are filled with random numbers and standardized to satisfy the general conditions. To make them further deviate from the orthogonal conditions, we intentionally choose $\boldsymbol{V}_{0,k}$ not orthogonal to $\boldsymbol{V}_k$ $(k = 1, 2)$, and artificially vary the norm of each column in $\boldsymbol{V}_{0,1}$ and $\boldsymbol{V}_{0,2}$.

- **Setting 3** (SIFA-B Model): The factors are generated in the same way as in Setting 2. The true loadings are specifically normalized to satisfy the orthogonal conditions. We note that the SIFA-B model is a special case of the SIFA-A model.

For each simulation setting, we fit JIVE, cov-JIVE, SIFA-A, and SIFA-B to the multiple simulated data sets, and fit PCA and SupSVD to the concatenated data $(\boldsymbol{Y}_1, \boldsymbol{Y}_2)$. We incorporate covariates for cov-JIVE, SIFA-A, SIFA-B, and SupSVD. In particular, when fitting the SIFA models, we set the functions in (10) to be linear, and use the ordinary least squares to estimate the coefficients. To avoid ambiguity, these model models are fitted with the true ranks. We set the rank for PCA and SupSVD to be $r_0 + r_1 + r_2$. We assess the performance of the LCV for rank selection in the next section.

To compare the loading estimation in JIVE, cov-JIVE, SIFA-A and SIFA-B, we use the Grassmannian metric (Mattila, 1999) between the true loadings in $\boldsymbol{V}_k$ and the estimated loadings in $\widehat{\boldsymbol{V}}_k$ for each $k = 0, 1, 2$. The metric is defined as $d_{\mathcal{G}}(\boldsymbol{V}_k, \widehat{\boldsymbol{V}}_k) = \sqrt{\sum_{i=1}^{r_k} \operatorname{acos}(\delta_i)^2}$, where δ_i is the ith singular value of $\boldsymbol{V}_k^T \widehat{\boldsymbol{V}}_k$. We also evaluate the maximal principal angle $\angle(\boldsymbol{V}, \widehat{\boldsymbol{V}})$ (Björck and Golub, 1973) between the subspaces in $\mathbb{R}^{p_1+p_2}$ spanned by the true loading vectors in $\boldsymbol{V} = (\boldsymbol{V}_0, \operatorname{blkdiag}(\boldsymbol{V}_1, \boldsymbol{V}_2))$ and the estimated ones, across all methods. To evaluate the accuracy of the estimated low-rank structure, we use $\|\boldsymbol{U}\boldsymbol{V}^T - \widehat{\boldsymbol{U}}\widehat{\boldsymbol{V}}^T\|_{\mathbb{F}}$ where $\boldsymbol{U} = (\boldsymbol{U}_0, \boldsymbol{U}_1, \boldsymbol{U}_2)$. The matrix $\widehat{\boldsymbol{U}}$ represents the estimated score matrix for PCA and JIVE, or the conditional expectation of the latent factor matrix for SupSVD, SIFA-A, and SIFA-B.

We also conduct comprehensive simulation studies to investigate: 1) the goodness of fit when the relations between covariates and latent factors are nonlinear; 2) the overfitting issue when nonparametric functions are used in the presence of linear relations; 3) the rank misspecification effect on the performance; 4) the violation of the Gaussian assumption; 5) the effect of rescaling different data sets; 6) the scalability

of SIFA-A and SIFA-B in high dimension. The simulation settings and results are contained in the supplementary material.

3.2 Rank Estimation by LCV

We briefly demonstrate the efficacy of the LCV method using a simulated example. Data are generated under Setting 3, with the chosen true ranks to be $(r_0, r_1, r_2) = (2, 3, 3)$. Additional examples under Setting 1 and Setting 2 are provided in the supplementary material. We particularly consider 9 candidate rank sets in the neighborhood of the true rank set: $(r_0, r_1, r_2) \in \{(1,2,2), (2,2,2), (3,2,2), (1,3,3), (2,3,3), (3,3,3), (3,4,3), (3,4,4), (4,4,4)\}$. We conduct a 10-fold LCV. The evaluated LCV scores (i.e., the negative log likelihood values of test samples) for different candidate sets in each cross validation run are shown in Figure 1. The average score reaches the minimum at the true rank set. Namely, the LCV method correctly selects the true ranks.

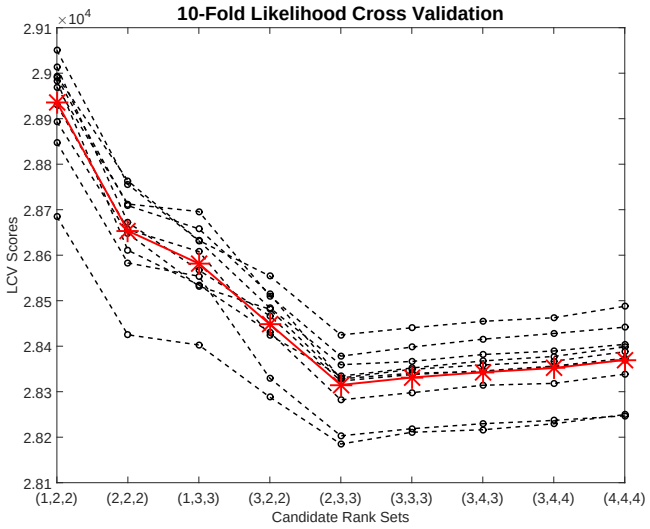

Figure 1: The LCV scores for 10-fold cross validation on 9 candidate rank sets. Each dashed line with circles contains corresponds to the LCV scores (negative log likelihood values) in one cross validation run. The solid line with stars contains the average LCV scores for different rank sets.

3.3 Simulation Results

For each setting, we repeat the simulation 100 times and summarize the results. The results are summarized in Table 1. In Setting 1, both SIFA-A and SIFA-B perform similarly to JIVE in terms of the loading estimation, even if the generative model is JIVE (i.e., the covariates are unrelated to the factors). Remarkably, the SIFA methods provide the best low-rank structure recovery accuracy among all. The reason is similar to the argument in Li et al. (2016): the shrinkage effect of the conditional expectation of the factors reduces estimation variance. In Setting 2, SIFA-A provides the best performance in all criteria. SIFA-B is suboptimal because the orthogonal conditions are severely violated. In Setting 3, the SIFA-B method performs the best, followed closely by SIFA-A. Both are significantly better than the competing methods. In practice, when the orthogonal conditions are well justified, SIFA-B is favorable due to the ultra-fast computation and accurate estimation. Otherwise, we recommend the use of the SIFA-A method.

4 GTEx Data Analysis

In this section, we apply the proposed method to the multi-tissue genetic data from the GTEx project. We use the phs000424.v6 data which are available at `http://www.gtexportal.org/` (registration required for data access). Technical details of data preprocessing and rank estimation can be found in the supplementary material.

The GTEx project collects gene expression data from multiple tissues (e.g., muscle, blood, skin) from the same set of subjects. We use the SIFA method to identify cross-tissue and tissue-specific gene expression patterns, and quantify the heritability of phenotypes representing expressions of a group of genes. Addressing the questions is integral to the fulfillment of the GTEx goal (Ardlie et al., 2015).

We particularly focus on the p53 signaling pathway in three tissues, i.e., muscle, blood, and skin, for the illustration purpose. The analysis can be easily generalized to other gene sets or tissues. After proper preprocessing and normalization, we obtain 191 genes on 204 common samples in each tissue, denoted by $\boldsymbol{Y}_1, \boldsymbol{Y}_2, \boldsymbol{Y}_3$. Each gene expression is standardized. In addition, we have the auxiliary data of sex, genotyping

Table 1: Simulation results under Setting 1, 2 and 3 (each with 100 simulation runs). The mean and standard deviation of each criterion for each method are shown in the table. The best results are highlighted in bold.

		SIFA-A	SIFA-B	JIVE	cov-JIVE	SupSVD	PCA
Setting 1 (JIVE)	$d_{\mathcal{G}}(\boldsymbol{V}_0, \widehat{\boldsymbol{V}}_0)$	0.61(0.03)	**0.60**(0.03)	0.69(0.07)	0.68(0.06)		
	$d_{\mathcal{G}}(\boldsymbol{V}_1, \widehat{\boldsymbol{V}}_1)$	0.82(0.06)	**0.81**(0.06)	0.91(0.17)	0.89(0.16)		
	$d_{\mathcal{G}}(\boldsymbol{V}_2, \widehat{\boldsymbol{V}}_2)$	1.32(0.17)	1.33(0.17)	1.31(0.18)	**1.30**(0.17)		
	$\angle(\boldsymbol{V}, \widehat{\boldsymbol{V}})$	64.73(10.98)	65.33(10.96)	65.16(11.31)	**64.35**(10.98)	87.02(2.74)	86.76(2.77)
	$\|\boldsymbol{U}\boldsymbol{V}^T - \widehat{\boldsymbol{U}}\widehat{\boldsymbol{V}}^T\|_{\mathbb{F}}$	193.28(2.85)	**193.26**(2.73)	240.49(4.53)	287.86(3.74)	239.21(2.82)	279.95(3.05)
Setting 2 (SIFA-A)	$d_{\mathcal{G}}(\boldsymbol{V}_0, \widehat{\boldsymbol{V}}_0)$	**0.37**(0.02)	1.01(0.05)	0.76(0.02)	1.40(0.14)		
	$d_{\mathcal{G}}(\boldsymbol{V}_1, \widehat{\boldsymbol{V}}_1)$	**0.27**(0.01)	1.06(0.08)	0.28(0.01)	1.41(0.03)		
	$d_{\mathcal{G}}(\boldsymbol{V}_2, \widehat{\boldsymbol{V}}_2)$	**0.52**(0.02)	0.67(0.03)	0.62(0.04)	1.80(0.09)		
	$\angle(\boldsymbol{V}, \widehat{\boldsymbol{V}})$	**27.67**(1.42)	44.46(2.42)	33.34(2.27)	88.30(1.43)	39.40(2.30)	46.97(4.28)
	$\|\boldsymbol{U}\boldsymbol{V}^T - \widehat{\boldsymbol{U}}\widehat{\boldsymbol{V}}^T\|_{\mathbb{F}}$	**169.21**(1.73)	207.74(2.03)	207.22(2.33)	296.14(2.11)	200.77(2.30)	235.44(3.00)
Setting 3 (SIFA-B)	$d_{\mathcal{G}}(\boldsymbol{V}_0, \widehat{\boldsymbol{V}}_0)$	0.38(0.03)	**0.30**(0.01)	0.61(0.02)	0.65(0.06)		
	$d_{\mathcal{G}}(\boldsymbol{V}_1, \widehat{\boldsymbol{V}}_1)$	0.25(0.01)	**0.24**(0.01)	0.25(0.01)	1.16(0.21)		
	$d_{\mathcal{G}}(\boldsymbol{V}_2, \widehat{\boldsymbol{V}}_2)$	0.35(0.01)	**0.34**(0.01)	0.36(0.01)	1.77(0.07)		
	$\angle(\boldsymbol{V}, \widehat{\boldsymbol{V}})$	15.03(0.56)	**14.97**(0.57)	15.86(0.73)	85.93(2.83)	26.21(1.09)	27.71(1.30)
	$\|\boldsymbol{U}\boldsymbol{V}^T - \widehat{\boldsymbol{U}}\widehat{\boldsymbol{V}}^T\|_{\mathbb{F}}$	171.99(1.65)	**171.51**(1.66)	204.64(1.97)	290.82(3.33)	200.77(1.81)	230.80(2.05)

platform index, and genetic variants for each sample as covariates. To reduce the dimensionality of the genetic variants, we obtain the top 30 principal components, which capture the majority of variation in the genotype data. The covariates are denoted by $\boldsymbol{X}$.

We first estimate the ranks for the joint and individual patterns. We use the two-step procedure described in Section 2.4, and exploit a variance explained criterion in the first step (with a preset 90% threshold). The joint and individual ranks are estimated to be $r_0 = 26$, $r_1 = 24$, $r_2 = 5$, and $r_3 = 20$. Note that the individual rank for blood ($r_2 = 5$) is much smaller than that for muscle or skin. From the viewpoint of the expression pattern richness, blood is very different from the other two tissues. This is generally concordant with the previous discoveries (Ardlie et al., 2015)

We fit a SIFA-B model to the data with linear relations between the covariates and the latent factors. For comparison, we also fit a JIVE model with the same ranks. The estimated joint and individual patterns are shown in Figure 2 (for the SIFA-B model) and Figure 3 (for the JIVE model). By taking into account the auxiliary covariates, the patterns obtained by the SIFA-B model are more discernable than those from the JIVE model. The joint structure in Figure 2 clearly captures the shared patterns among samples across tissues, while the individual structure distinguishes different tissues. We also quantify the variation explained by different parts in both methods (see Figure 4). The SIFA-B decomposition attributes more variation to the individual structure than the JIVE method, which is consistent with the domain knowledge that the p53 gene expressions are highly tissue specific (Ribeiro et al., 2001; Tendler et al., 1999). The tissue-specific expression patterns may be used to investigate tissue identity and functions.

To quantify the heritability of the derived phenotypes (i.e., joint and individual scores) representing the p53 gene expressions, we calculate the variation explained by different components of the SIFA model. The results are summarized in Table 2. Within the joint structure (common across all tissues), the genetic variants explain about 17% of the variation, which is concordant with the general belief in the literature (Brown et al., 2015). The sex and the platform information take up 2% and 2.5% of the variation, respectively. The vast majority of the variation remains unexplained, which provokes further investigation. The individual structure for each tissue has a

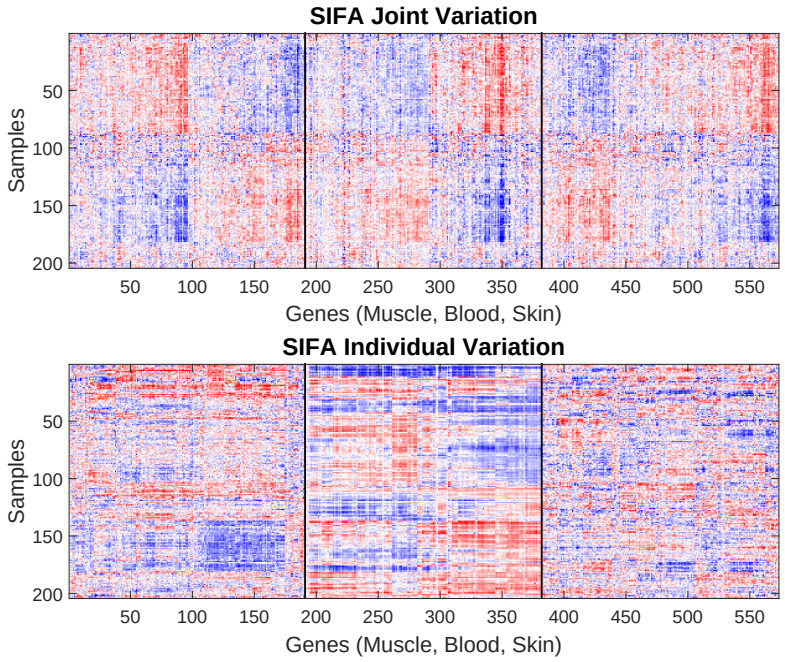

Figure 2: GTEx Example: The heat maps of the joint and individual gene expression patterns for the p53 signaling pathway identified by the SIFA-B model. For visualization purpose, we reorder samples across three tissues and genes in each tissue. Top panel: the joint structure in three tissues; Bottom panel: the individual structures in three tissues. In each panel, the samples are matched across tissues.

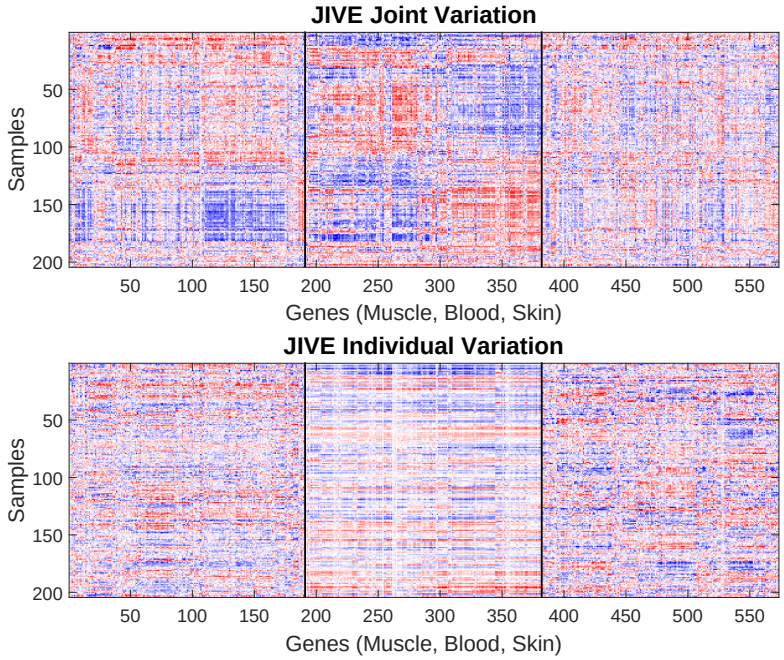

Figure 3: GTEx Example: The heat maps of the joint and individual gene expression patterns for the p53 signaling pathway identified by the JIVE model. The rows and columns are ordered in the same way as in Figure 2.

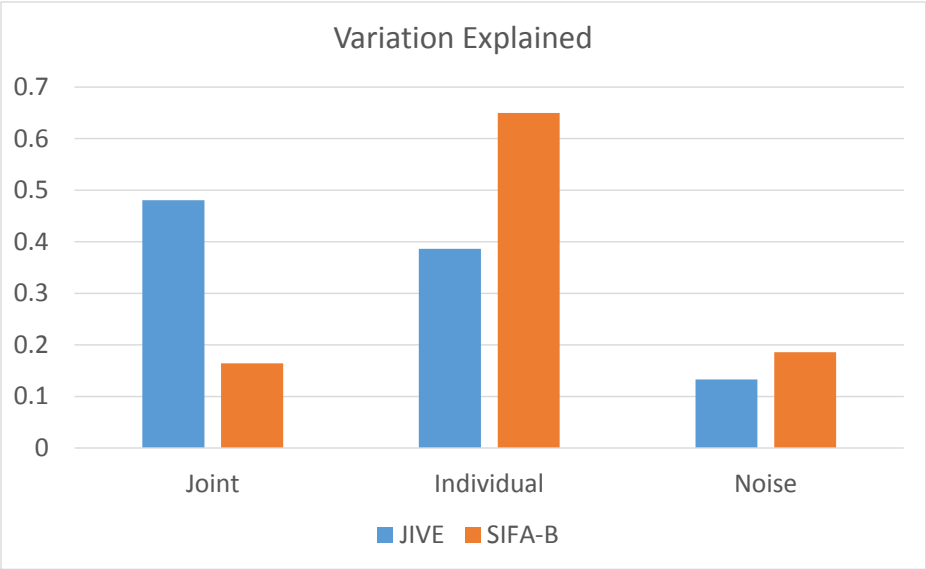

Figure 4: GTEx Example: The variation explained by different components in the JIVE model and the SIFA-B model, respectively.

similar decomposition to the joint structure. An interesting finding is that sex is not a major contributor to the individual gene expression patterns in blood. The derived pathway expression phenotypes could also potentially be used to discover associations with clinical outcome and environmental factors. Due to the lack of such information in the GTEx data, we do not further pursue it here.

5 Discussion

In this paper, we develop a supervised integrated factor analysis framework for reduction and integration of multi-view data. It decomposes multiple related data sets into joint and individual structures, while incorporating covariate supervision through parametric or nonparametric models. We investigate the identifiability of the model under two sets of conditions, the general conditions and the orthogonal conditions, each being useful in separate situations. An efficient EM algorithm with some variants is devised to fit the model. In particular, it is very easy to capture nonlinear relations between covariates and latent factors, and to incorporate variable selection of covariates. The comprehensive simulation studies demonstrate the efficacy of the

Table 2: GTEx Example: The genetic variation explained by different factors in different tissues. For each tissue, the last column gives the percentage explained by the joint and individual structure, and the noise (add up to 1). The variation in the joint (individual) structure is further attributed to the genotype, sex, platform, and other unknown sources (add up to 1).

		genotype	Sex	Platform	Unknown	Total
Muscle	Joint	17.09%	2.03%	2.56%	78.32%	16.44%
	Individual	15.55%	2.66%	1.74%	80.05%	65.29%
	Noise					18.27%
Blood	Joint	17.09%	2.03%	2.56%	78.32%	16.44%
	Individual	14.05%	0.65%	0.90%	84.39%	63.56%
	Noise					20.00%
Skin	Joint	17.09%	2.03%	2.56%	78.32%	16.44%
	Individual	16.55%	1.52%	1.04%	80.89%	66.06%
	Noise					17.50%

proposed methods. With application to the GTEx data, we provide new insights into the genetic variation of a gene set across multiple tissues.

There are several directions for future research. First of all, it is of potential interest to generalize the current framework to accommodate non-normal data. Second, the model may be specially modified to capture partially joint structure pertaining to multiple but not all data sets. This is especially relevant when multiple data sets are naturally grouped at the source level. Third, customized rank estimation methods need further investigation.

Acknowledgements

Gen Li was partially supported by the Calderone Junior Faculty Award by the Mailman School of Public Health at Columbia University. Sungkyu Jung was partially supported by the National Science Foundation grant DMS-1307178.

References

Abdi, H., Williams, L. J., and Valentin, D. (2013). Multiple factor analysis: principal component analysis for multitable and multiblock data sets. *Wiley Interdisciplinary reviews: computational statistics* **5,** 149–179.

Ahn, J. and Marron, J. (2010). The maximal data piling direction for discrimination. *Biometrika* **97,** 254–259.

Ardlie, K. G., Deluca, D. S., et al. (2015). The genotype-tissue expression (gtex) pilot analysis: Multitissue gene regulation in humans. *Science* **348,** 648–660.

Björck, k. and Golub, G. H. (1973). Numerical methods for computing angles between linear subspaces. *Mathematics of computation* **27,** 579–594.

Brown, A., Ding, Z., Viñuela, A., Glass, D., Parts, L., Spector, T., Winn, J., and Durbin, R. (2015). Pathway based factor analysis of gene expression data produces highly heritable phenotypes that associate with age. *G3: Genes— Genomes— Genetics* pages g3–114.

Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., et al. (2004). Least angle regression. *The Annals of statistics* **32,** 407–499.

Fan, J. and Gijbels, I. (1996). *Local polynomial modelling and its applications: monographs on statistics and applied probability*, volume 66. CRC Press.

Fan, J., Liao, Y., and Wang, W. (2016). Projected principal component analysis in factor models. *Annals of statistics* **44,** 219–254.

Gower, J. C. and Dijksterhuis, G. B. (2004). *Procrustes problems*, volume 3. Oxford University Press Oxford.

Harville, D. A. (1997). *Matrix Algebra from a Statistician's Perspective*. Springer New York.

Hastie, T. J. and Tibshirani, R. J. (1990). *Generalized additive models*, volume 43. CRC Press.

Hollander, M., Wolfe, D. A., and Chicken, E. (2013). *Nonparametric statistical methods*. John Wiley & Sons.

Lafferty, J. and Wasserman, L. (2008). Rodeo: sparse, greedy nonparametric regression. *The Annals of Statistics* **36,** 28–63.

Lee, S. and Jung, S. (2016). Combined analysis of amplitude and phase variations in functional data. *arXiv preprint arXiv:1603.01775* .

Li, G., Yang, D., Nobel, A. B., and Shen, H. (2016). Supervised singular value decomposition and its asymptotic properties. *Journal of Multivariate Analysis* **146,** 7–17.

Lock, E. F., Hoadley, K. A., Marron, J. S., and Nobel, A. B. (2013). Joint and individual variation explained (JIVE) for integrated analysis of multiple data types. *Annals of Applied Statistics* **7,** 523–542.

Löfstedt, T., Hoffman, D., and Trygg, J. (2013). Global, local and unique decompositions in onpls for multiblock data analysis. *Analytica chimica acta* **791,** 13–24.

Mattila, P. (1999). *Geometry of sets and measures in Euclidean spaces: fractals and rectifiability*, volume 44. Cambridge University Press.

Ramsay, J. O. and Silverman, B. W. (2002). *Applied functional data analysis: methods and case studies*, volume 77. Citeseer.

Ravikumar, P., Lafferty, J., Liu, H., and Wasserman, L. (2009). Sparse additive models. *Journal of the Royal Statistical Society: Series B* **71,** 1009–1030.

Ray, P., Zheng, L., Lucas, J., and Carin, L. (2014). Bayesian joint analysis of heterogeneous genomics data. *Bioinformatics* **30,** 1370–1376.

Ribeiro, R. C., Sandrini, F., et al. (2001). An inherited p53 mutation that contributes in a tissue-specific manner to pediatric adrenal cortical carcinoma. *Proceedings of the National Academy of Sciences* **98,** 9330–9335.

Schouteden, M., Van Deun, K., Wilderjans, T. F., and Van Mechelen, I. (2014). Performing disco-sca to search for distinctive and common information in linked data. *Behavior research methods* **46,** 576–587.

Shabalin, A. and Nobel, A. (2013). Reconstruction of a low-rank matrix in the presence of Gaussian noise. *Journal of Multivariate Analysis* **118,** 67–76.

Tendler, Y., Weisinger, G., et al. (1999). Tissue-specific p53 expression in the nervous system. *Molecular brain research* **72,** 40–46.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B* **58,** 267–288.

Tipping, M. E. and Bishop, C. M. (1999). Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B* **61,** 611–622.

Tseng, G. C., Ghosh, D., and Zhou, X. J. (2015). *Integrating Omics Data*. Cambridge University Press.

Yang, Z. and Michailidis, G. (2016). A non-negative matrix factorization method for detecting modules in heterogeneous omics multi-modal data. *Bioinformatics* **32,** 1–8.

Zhou, G., Cichocki, A., Zhang, Y., and Mandic, D. P. (2016). Group component analysis for multiblock data: Common and individual feature extraction. *IEEE Transactions on Neural Networks and Learning Systems* **27,** 2426–2439.

Supplementary Materials for "Incorporating Covariates into Integrated Factor Analysis of Multi-View Data" by Gen Li, and Sungkyu Jung

A Identifiability

In this section, we first give a few examples of unidentifiable models without any additional conditions other than the basic conditions (i.e., $\boldsymbol{V}_k^T\boldsymbol{V}_k = \boldsymbol{I}$ and $\boldsymbol{\Sigma}_k$ being diagonal with positive, distinct and decreasing eigenvalues for each $k = 0, 1, \cdots, K$). Then we prove Proposition 2.1, the identifiability of the SIFA model under the general conditions.

A.1 Examples of Unidentifiable Models

By inspecting the log likelihood function (7) of the SIFA model, we can see that any two sets of model parameters are equivalent as long as they lead to the equal values of $\boldsymbol{\mu}_\star$ and $\boldsymbol{\Sigma}_\star$, where

$$\boldsymbol{\mu}_\star = [\boldsymbol{f}_0(\boldsymbol{X}), \boldsymbol{f}_1(\boldsymbol{X}), \cdots, \boldsymbol{f}_K(\boldsymbol{X})](\boldsymbol{V}_0, \boldsymbol{V}_\star)^T; \tag{S.1}$$

$$\boldsymbol{\Sigma}_\star = \boldsymbol{V}_0\boldsymbol{\Sigma}_0\boldsymbol{V}_0^T + \boldsymbol{V}_\star\boldsymbol{\Sigma}_{\boldsymbol{F}}\boldsymbol{V}_\star^T + \boldsymbol{\Sigma}_{\boldsymbol{E}}. \tag{S.2}$$

In the following, we provide two examples of equivalent models. One example involves models with different ranks, and the other example concerns models with the same set of ranks.

Example 1 (models with different ranks): Consider a very simple example with two primary data sets ($K = 2$). Suppose one model parameter set $\theta = \{\boldsymbol{f}_0(\cdot), \boldsymbol{f}_1(\cdot), \boldsymbol{f}_2(\cdot), \boldsymbol{V}_0, \boldsymbol{V}_1, \boldsymbol{V}_2, \boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_1, \boldsymbol{\Sigma}_2,$
$\sigma_1^2, \sigma_2^2\}$ satisfies the basic conditions and has rank $r_0 = 1, r_1 = r_2 = 2$. Note that the combined loading matrix $(\boldsymbol{V}_0, \boldsymbol{V}_\star)$, where $\boldsymbol{V}_\star = \text{blkdiag}(\boldsymbol{V}_1, \boldsymbol{V}_2)$, may not have orthonormal columns.

One idea of generating an equivalent parameter set $\widehat{\theta}$ to θ is to intentionally mistake an individual component to a joint component and re-standardize everything

accordingly. We define

$$\widetilde{\boldsymbol{V}_0} = \left(\boldsymbol{V}_0, \begin{pmatrix} \boldsymbol{v}_{1,1} \\ \mathbf{0} \end{pmatrix}\right), \widehat{\boldsymbol{V}_1} = \boldsymbol{v}_{1,2}, \widetilde{\boldsymbol{\Sigma}_0} = \begin{pmatrix} \boldsymbol{\Sigma}_0 & \mathbf{0} \\ \mathbf{0} & \Sigma_{1,1} \end{pmatrix}, \widehat{\boldsymbol{\Sigma}_1} = \Sigma_{1,2}, \widetilde{\boldsymbol{f}_0} = (\boldsymbol{f}_0, f_{1,1}), \widehat{\boldsymbol{f}_1} = f_{1,2},$$

where $\boldsymbol{v}_{1,1}$ and $\boldsymbol{v}_{1,2}$ are the first and second columns of $\boldsymbol{V}_1$, $\Sigma_{1,1}$ and $\Sigma_{1,2}$ are the first and second diagonal values of $\boldsymbol{\Sigma}_1$, and $f_{1,1}$ and $f_{1,2}$ are the first and second components of $\boldsymbol{f}_1$. In addition, let $\widehat{\boldsymbol{V}_2} = \boldsymbol{V}_2$, $\widehat{\boldsymbol{\Sigma}_2} = \boldsymbol{\Sigma}_2$, $\widehat{\boldsymbol{f}_2} = \boldsymbol{f}_2$, $\widehat{\sigma_1^2} = \sigma_1^2$, and $\widehat{\sigma_2^2} = \sigma_2^2$.

The matrix $\widetilde{\boldsymbol{V}_0}$ is an intermediate loading matrix which may not have orthonormal columns (thus violating the basic conditions). A new loading matrix $\widehat{\boldsymbol{V}_0}$ and a new covariance matrix $\widehat{\boldsymbol{\Sigma}_0}$ are obtained by the eigendecomposition of $\widetilde{\boldsymbol{V}_0}\widetilde{\boldsymbol{\Sigma}_0}\widetilde{\boldsymbol{V}_0}^T$, where $\widehat{\boldsymbol{V}_0}$ contains the eigenvectors and $\widehat{\boldsymbol{\Sigma}_0}$ contains the eigenvalues. As a result, $\widehat{\boldsymbol{V}_0}$ and $\widehat{\boldsymbol{\Sigma}_0}$ satisfy the basic conditions and (S.2) remains unchanged. Since the column spaces of $\widetilde{\boldsymbol{V}_0}$ and $\widehat{\boldsymbol{V}_0}$ are the same, we can easily obtain a 3×3 matrix $\boldsymbol{Q}$ such that $\widetilde{\boldsymbol{V}_0} = \widehat{\boldsymbol{V}_0}\boldsymbol{Q}$. Correspondingly, we set $\widehat{\boldsymbol{f}_0}(\boldsymbol{X}) = \widetilde{\boldsymbol{f}_0}(\boldsymbol{X})\boldsymbol{Q}$.

As a result, we get a new parameter set $\widehat{\theta} = \{\widehat{\boldsymbol{f}_0}(\cdot), \widehat{\boldsymbol{f}_1}(\cdot), \widehat{\boldsymbol{f}_2}(\cdot), \widehat{\boldsymbol{V}_0}, \widehat{\boldsymbol{V}_1}, \widehat{\boldsymbol{V}_2}, \widehat{\boldsymbol{\Sigma}_0}, \widehat{\boldsymbol{\Sigma}_1}, \widehat{\boldsymbol{\Sigma}_2}, \widehat{\sigma_1^2}, \widehat{\sigma_2^2}\}$. It is easy to check that $\widehat{\boldsymbol{V}_k}^T\widehat{\boldsymbol{V}_k} = \boldsymbol{I}$ and $\widehat{\boldsymbol{\Sigma}_k}$'s are diagonal for $k = 0, 1, 2$ (with a little care we can find an example where the resulting $\widehat{\boldsymbol{\Sigma}_k}$ have distinct eigenvalues). Namely, $\widehat{\theta}$ is a qualified parameter set with $r_0 = 2, r_1 = 1, r_2 = 2$. According to the construction process, it is also trivial to see that θ and $\widehat{\theta}$ give exactly the same values in (S.1) and (S.2). Therefore, θ and $\widehat{\theta}$ are equivalent but distinct model parameter sets.

Example 2 (models with the same set of ranks): Let us consider the same model parameter set θ as in Example 1. Instead of padding the first column of $\boldsymbol{V}_1$ with zeros and making it a joint loading vector, we pad the second column and make it a joint loading vector. Specifically, let

$$\widetilde{\boldsymbol{V}_0}' = \left(\boldsymbol{V}_0, \begin{pmatrix} \boldsymbol{v}_{1,2} \\ \mathbf{0} \end{pmatrix}\right), \widehat{\boldsymbol{V}_1}' = \boldsymbol{v}_{1,1}, \widetilde{\boldsymbol{\Sigma}_0}' = \begin{pmatrix} \boldsymbol{\Sigma}_0 & \mathbf{0} \\ \mathbf{0} & \Sigma_{1,2} \end{pmatrix}, \widehat{\boldsymbol{\Sigma}_1}' = \Sigma_{1,1}, \widetilde{\boldsymbol{f}_0}' = (\boldsymbol{f}_0, f_{1,2}), \widehat{\boldsymbol{f}_1}' = f_{1,1},$$

and $\widehat{\boldsymbol{V}_2}' = \boldsymbol{V}_2$, $\widehat{\boldsymbol{\Sigma}_2}' = \boldsymbol{\Sigma}_2$, $\widehat{\boldsymbol{f}_2}' = \boldsymbol{f}_2$, $\widehat{\sigma_1^2}' = \sigma_1^2$, and $\widehat{\sigma_2^2}' = \sigma_2^2$. Following exactly the same procedure as in Example 1, we can construct another equivalent parameter set $\widehat{\theta}'$ to θ, with the ranks $r_0 = 2, r_1 = 1, r_2 = 2$. Due to the transitive relation in the

equivalent set, $\widehat{\theta}$ is equivalent to $\widehat{\theta}'$, and both parameter sets have the same set of ranks. We note that $\widehat{\theta}'$ is not identical to $\widehat{\theta}$ because $\widehat{\boldsymbol{V}_1} = \boldsymbol{v}_{1,2} \neq \boldsymbol{v}_{1,1} = \widehat{\boldsymbol{V}_1}'$. As a result, we obtain two equivalent but distinct parameter sets with the same set of ranks.

In both examples, the parameter sets $\widehat{\theta}$ and $\widehat{\theta}'$ do not satisfy the general conditions. Next, we shall prove that under the general conditions, the parameter set is unique.

A.2 Proof of Proposition 2.1

Proof. Suppose $\theta = \{\boldsymbol{f}_0(\cdot), \boldsymbol{f}_k(\cdot), \boldsymbol{V}_0, \boldsymbol{V}_k, \boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_k, \sigma_k^2;\ k = 1, \cdots, K\}$ and $\widehat{\theta} = \{\widehat{\boldsymbol{f}_0}(\cdot), \widehat{\boldsymbol{f}_k}(\cdot), \widehat{\boldsymbol{V}_0}, \widehat{\boldsymbol{V}_k}, \widehat{\boldsymbol{\Sigma}_0}, \widehat{\boldsymbol{\Sigma}_k}, \widehat{\sigma_k^2};\ k = 1, \cdots, K\}$ are two equivalent model parameter sets that satisfy the basic conditions and the general conditions A1 and A2. In the following, we will first prove $\widehat{\sigma_k^2} = \sigma_k^2$ for $k = 1, \cdots, K$. Then we will prove that both parameter sets have the same set of ranks, and the corresponding joint rank r_0 is minimal in the equivalent class. Next, for each $k = 0, 1, \cdots, K$, we will prove $\widehat{\boldsymbol{V}_k} = \boldsymbol{V}_k$ (up to trivial column-wise sign changes) and $\widehat{\boldsymbol{\Sigma}_k} = \boldsymbol{\Sigma}_k$. Finally, we will prove $\widehat{\boldsymbol{f}_k} = \boldsymbol{f}_k$ (again, up to trivial sign changes).

Equal Variance: Let $\boldsymbol{\Sigma}_\star$ and $\widehat{\boldsymbol{\Sigma}_\star}$ denote the covariance matrices in (S.2) derived from the respective parameter sets. Since $\boldsymbol{\Sigma}_\star = \widehat{\boldsymbol{\Sigma}_\star}$, all corresponding submatrices are equal. In particular, we focus on the first $p_1 \times p_1$ submatrix, and have $\boldsymbol{V}_{0,1}\boldsymbol{\Sigma}_0\boldsymbol{V}_{0,1}^T + \boldsymbol{V}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T + \sigma_1^2\boldsymbol{I} = \widehat{\boldsymbol{V}_{0,1}}\widehat{\boldsymbol{\Sigma}_0}\widehat{\boldsymbol{V}_{0,1}}^T + \widehat{\boldsymbol{V}_1}\widehat{\boldsymbol{\Sigma}_1}\widehat{\boldsymbol{V}_1}^T + \widehat{\sigma_1^2}\boldsymbol{I}$, or equivalently,

$$(\boldsymbol{V}_{0,1}, \boldsymbol{V}_1)\begin{pmatrix}\boldsymbol{\Sigma}_0 & \boldsymbol{0}\\ \boldsymbol{0} & \boldsymbol{\Sigma}_1\end{pmatrix}\begin{pmatrix}\boldsymbol{V}_{0,1}^T\\ \boldsymbol{V}_1^T\end{pmatrix} + \sigma_1^2\boldsymbol{I} = (\widehat{\boldsymbol{V}_{0,1}}, \widehat{\boldsymbol{V}_1})\begin{pmatrix}\widehat{\boldsymbol{\Sigma}_0} & \boldsymbol{0}\\ \boldsymbol{0} & \widehat{\boldsymbol{\Sigma}_1}\end{pmatrix}\begin{pmatrix}\widehat{\boldsymbol{V}_{0,1}}^T\\ \widehat{\boldsymbol{V}_1}^T\end{pmatrix} + \widehat{\sigma_1^2}\boldsymbol{I}. \tag{S.3}$$

From Condition A2, we know $r_0 + r_1 < p_1$ and $\widehat{r_0} + \widehat{r_1} < p_1$, where r_k and $\widehat{r_k}$ are the numbers of columns in $\boldsymbol{V}_k$ and $\widehat{\boldsymbol{V}_k}$ respectively ($k = 0, 1 \cdots, K$). Thus, the first term on both sides of the equation is a low-rank non-negative definite matrix. It is easy to see the smallest eigenvalue of the left-hand side (LHS) matrix is σ_1^2 and the smallest eigenvalue of the right-hand side (RHS) matrix is $\widehat{\sigma_1^2}$. Therefore, we have $\sigma_1^2 = \widehat{\sigma_1^2}$. Similar argument for the kth diagonal submatrix leads to $\sigma_k^2 = \widehat{\sigma_k^2},\ k = 1, \cdots, K$.

Equal Ranks: Taking out the second term from both sides of (S.3), we have

$$(\boldsymbol{V}_{0,1}, \boldsymbol{V}_1) \begin{pmatrix} \boldsymbol{\Sigma}_0 & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_1 \end{pmatrix} \begin{pmatrix} \boldsymbol{V}_{0,1}^T \\ \boldsymbol{V}_1^T \end{pmatrix} = (\widehat{\boldsymbol{V}_{0,1}}, \widehat{\boldsymbol{V}_1}) \begin{pmatrix} \widehat{\boldsymbol{\Sigma}_0} & \mathbf{0} \\ \mathbf{0} & \widehat{\boldsymbol{\Sigma}_1} \end{pmatrix} \begin{pmatrix} \widehat{\boldsymbol{V}_{0,1}}^T \\ \widehat{\boldsymbol{V}_1}^T \end{pmatrix}.$$

Since $(\boldsymbol{V}_{0,1}, \boldsymbol{V}_1)$ and $(\widehat{\boldsymbol{V}_{0,1}}, \widehat{\boldsymbol{V}_1})$ both have full column ranks according to Condition A1 and A2, and the diagonal values of blkdiag$(\boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_1)$ and blkdiag$(\widehat{\boldsymbol{\Sigma}_0}, \widehat{\boldsymbol{\Sigma}_1})$ are positive, the LHS matrix and the RHS matrix have ranks $r_0 + r_1$ and $\widehat{r_0} + \widehat{r_1}$, respectively. Consequently, we get k equations: $r_0 + r_k = \widehat{r_0} + \widehat{r_k}$ for $k = 1, \cdots, K$. In addition, from $\boldsymbol{\Sigma}_\star = \widehat{\boldsymbol{\Sigma}_\star}$, we get another equation: $r_0 + \sum_{k=1}^K r_k = \widehat{r_0} + \sum_{k=1}^K \widehat{r_k}$. Solving these $K+1$ equations, we get $r_k = \widehat{r_k}$, for $k = 0, 1, \cdots, K$.

To see that the r_0 of θ is actually minimal in the equivalent class (in which parameter sets do not necessarily satisfy the general conditions), we can look at an off-diagonal block matrix of $\boldsymbol{\Sigma}_\star$. For example, the submatrix consisting of the (p_1+1)th to the (p_1+p_2)th columns and the first to the p_1th rows is $\boldsymbol{V}_{0,1}\boldsymbol{\Sigma}_0\boldsymbol{V}_{0,2}^T$. It has rank r_0. For any parameter set in the equivalent class, the derived submatrix should be identical with $\boldsymbol{V}_{0,1}\boldsymbol{\Sigma}_0\boldsymbol{V}_{0,2}^T$ and thus also have rank r_0. Consequently, it must have no fewer than r_0 joint loadings. Namely, r_0 is minimal in the equivalent class.

Equal Loading and Covariance: In order to prove the loading matrices are equal, we first look at the off-diagonal block matrices of $\boldsymbol{\Sigma}_\star$ and $\widehat{\boldsymbol{\Sigma}_\star}$. In particular, we have

$$\boldsymbol{V}_{0,k_1}\boldsymbol{\Sigma}_0\boldsymbol{V}_{0,k_2}^T = \widehat{\boldsymbol{V}_{0,k_1}}\widehat{\boldsymbol{\Sigma}_0}\widehat{\boldsymbol{V}_{0,k_2}}^T, \ k_1 \neq k_2 \in \{1, \cdots, K\}. \tag{S.4}$$

Based on Condition A1, the LHS and RHS matrices both have rank r_0, and the column spaces are col$(\boldsymbol{V}_{0,k_1})$ and col$(\widehat{\boldsymbol{V}_{0,k_1}})$ respectively. Thus we have col$(\boldsymbol{V}_{0,k})$ = col$(\widehat{\boldsymbol{V}_{0,k}})$ for $k = 1, \cdots, K$.

Then we look at the diagonal block matrices. In particular, we have

$$\boldsymbol{V}_{0,1}\boldsymbol{\Sigma}_0\boldsymbol{V}_{0,1}^T + \boldsymbol{V}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T = \widehat{\boldsymbol{V}_{0,1}}\widehat{\boldsymbol{\Sigma}_0}\widehat{\boldsymbol{V}_{0,1}}^T + \widehat{\boldsymbol{V}_1}\widehat{\boldsymbol{\Sigma}_1}\widehat{\boldsymbol{V}_1}^T. \tag{S.5}$$

Let $\boldsymbol{C} = \boldsymbol{V}_{0,1}\boldsymbol{\Sigma}_0\boldsymbol{V}_{0,1}^T - \widehat{\boldsymbol{V}_{0,1}}\widehat{\boldsymbol{\Sigma}_0}\widehat{\boldsymbol{V}_{0,1}}^T$. From the previous discussion, we know col$(\boldsymbol{C}) \subseteq$ col$(\boldsymbol{V}_{0,1})$. Plugging $\boldsymbol{C}$ into (S.5), we have

$$\widehat{\boldsymbol{V}_1}\widehat{\boldsymbol{\Sigma}_1}\widehat{\boldsymbol{V}_1}^T = \boldsymbol{V}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T + \boldsymbol{C}.$$

According to Corollary 18.5.5 in Harville (1997), the rank of the RHS matrix is bounded below by $\text{rank}(\boldsymbol{V}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T) + \text{rank}(\boldsymbol{C}) - 2c$, where c is the dimension of $\text{col}(\boldsymbol{V}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T) \cap \text{col}(\boldsymbol{C})$. Since $\text{col}(\boldsymbol{V}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T) = \text{col}(\boldsymbol{V}_1)$, $\text{col}(\boldsymbol{C}) \subseteq \text{col}(\boldsymbol{V}_{0,1})$, and $\text{col}(\boldsymbol{V}_1) \cap \text{col}(\boldsymbol{V}_{0,1}) = \{\boldsymbol{0}\}$ (according to Condition A2), we know $c = 0$. Therefore, we have

$$\text{rank}(\widehat{\boldsymbol{V}_1}\widehat{\boldsymbol{\Sigma}_1}\widehat{\boldsymbol{V}_1}^T) = \text{rank}(\boldsymbol{V}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T + \boldsymbol{C}) \geq \text{rank}(\boldsymbol{V}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T) + \text{rank}(\boldsymbol{C}) = r_1 + \text{rank}(\boldsymbol{C}).$$

In combination with the fact that $\text{rank}(\widehat{\boldsymbol{V}_1}\widehat{\boldsymbol{\Sigma}_1}\widehat{\boldsymbol{V}_1}^T) = \widehat{r_1} = r_1$, we get $\text{rank}(\boldsymbol{C}) = 0$, i.e., $\boldsymbol{C} = \boldsymbol{0}$.

Immediately, we have $\widehat{\boldsymbol{V}_k}\widehat{\boldsymbol{\Sigma}_k}\widehat{\boldsymbol{V}_k}^T = \boldsymbol{V}_k\boldsymbol{\Sigma}_k\boldsymbol{V}_k^T$ and $\widehat{\boldsymbol{V}_{0,k}}\widehat{\boldsymbol{\Sigma}_0}\widehat{\boldsymbol{V}_{0,k}}^T = \boldsymbol{V}_{0,k}\boldsymbol{\Sigma}_0\boldsymbol{V}_{0,k}^T$ for $k = 1, \cdots, K$. Combining with (S.4), we also have $\widehat{\boldsymbol{V}_0}\widehat{\boldsymbol{\Sigma}_0}\widehat{\boldsymbol{V}_0}^T = \boldsymbol{V}_0\boldsymbol{\Sigma}_0\boldsymbol{V}_0^T$. Under the basic conditions, $\boldsymbol{V}_k^T\boldsymbol{\Sigma}_k\boldsymbol{V}_k$ and $\widehat{\boldsymbol{V}_k}^T\widehat{\boldsymbol{\Sigma}_k}\widehat{\boldsymbol{V}_k}$ $(k = 0, 1, \cdots, K)$ are both in the form of eigendecomposition. From the uniqueness of eigendecomposition, we conclude that

$$\boldsymbol{V}_k = \widehat{\boldsymbol{V}_k}\boldsymbol{S}_k,\ \boldsymbol{\Sigma}_k = \widehat{\boldsymbol{\Sigma}_k},\ k = 0, 1, \cdots, K,$$

where $\boldsymbol{S}_k$ is a diagonal matrix (with dimension compatible with $\widehat{\boldsymbol{V}_k}$) whose diagonal values are either 1 or -1. For an eigenvector, the sign change is trivial. As discussed in the main article, one could easily fix the sign of each eigenvector by setting the first nonzero entry to be positive. Without loss of generality, we assume $\boldsymbol{S}_k$ is the identity matrix.

Equal Functions: We plug in $\boldsymbol{V}_0 = \widehat{\boldsymbol{V}_0}$ and $\boldsymbol{V}_\star = \widehat{\boldsymbol{V}_\star}$ in (S.1). As a result, we obtain that $\boldsymbol{f}_k(\boldsymbol{X}) - \widehat{\boldsymbol{f}_k}(\boldsymbol{X}) = \boldsymbol{0}$ holds for any $\boldsymbol{X}$ $(k = 0, 1, \cdots, K)$. Namely, $\boldsymbol{f}_k = \widehat{\boldsymbol{f}_k}$, $k = 0, 1, \cdots, K$. □

B Algorithm Details

In this section, we provide a detailed description of the EM algorithm for fitting an SIFA model. The M steps for different sets of identifiability conditions are slightly different, while the E step is universal. In the following, we will derive the conditional distribution of the latent variables, and then derive the specifics of the M steps under different conditions. We will use the notations used in the main article wherever applicable.

B.1 E Step

Assume $\theta^{(l)}$ is the parameter set estimated from the lth iteration. We focus on the estimation in the $(l+1)$th iteration. For simplicity, we shall drop the superscript when it does not cause any confusion. Under the normal assumption, the conditional distribution of $\boldsymbol{y}_{\star(i)}$ (a column vector of the ith row of $\boldsymbol{Y}_\star$) given $\boldsymbol{u}_{0(i)}$ and $\boldsymbol{u}_{\star(i)}$ (column vectors of the ith row of $\boldsymbol{U}_0$ and $\boldsymbol{U}_\star$) is

$$\boldsymbol{y}_{\star(i)}|\boldsymbol{u}_{0(i)}, \boldsymbol{u}_{\star(i)} \ \sim \ \mathcal{N}(\boldsymbol{\mu}_{(i)}, \boldsymbol{\Sigma_E}),$$

where $\boldsymbol{\mu}_{(i)}$ is a column vector of the ith row of $\boldsymbol{U}_0\boldsymbol{V}_0^T + \boldsymbol{U}_\star\boldsymbol{V}_\star^T$. The marginal distributions of $\boldsymbol{u}_{0(i)}$ and $\boldsymbol{u}_{\star(i)}$ are

$$\begin{aligned}
\boldsymbol{u}_{0(i)} &\sim \mathcal{N}\left(\boldsymbol{f}_0(\boldsymbol{x}_{(i)})^T, \boldsymbol{\Sigma}_0\right), \\
\boldsymbol{u}_{\star(i)} &\sim \mathcal{N}\left([\boldsymbol{f}_1(\boldsymbol{x}_{(i)}), \cdots, \boldsymbol{f}_K(\boldsymbol{x}_{(i)})]^T, \boldsymbol{\Sigma_F}\right),
\end{aligned}$$

where $\boldsymbol{f}_k(\boldsymbol{x}_{(i)})$ is the ith row of $\boldsymbol{f}_k(\boldsymbol{X})$ for $k = 0, 1, \cdots, K$. With some basic algebraic calculation, we can easily get the marginal distribution of $\boldsymbol{y}_{\star(i)}$ and the joint distribution of $\boldsymbol{y}_{\star(i)}, \boldsymbol{u}_{0(i)}, \boldsymbol{u}_{\star(i)}$ as a normal distribution. To avoid overcomplicated notations, we will not present the intermediate results here.

Subsequently, we obtain the conditional distribution of $\boldsymbol{u}_{0(i)}, \boldsymbol{u}_{\star(i)}$ given $\boldsymbol{y}_{\star(i)}$, which is a normal distribution. In particular, the matrix form of the conditional mean is

$$\mathbb{E}(\boldsymbol{U}_0, \boldsymbol{U}_\star|\boldsymbol{Y}_\star) = [\boldsymbol{f}_0(\boldsymbol{X}), \boldsymbol{f}_1(\boldsymbol{X}), \cdots, \boldsymbol{f}_K(\boldsymbol{X})] + (\boldsymbol{Y}_\star - \boldsymbol{\mu}_\star)\boldsymbol{\Sigma}_\star^{-1}(\boldsymbol{V}_0\boldsymbol{\Sigma}_0, \boldsymbol{V}_\star\boldsymbol{\Sigma_F}) \quad \text{(S.6)}$$

where $\boldsymbol{\Sigma}_\star$ is the marginal covariance matrix of $\boldsymbol{y}_{\star(i)}$, which has the form (S.2). The conditional covariance matrix of each row of $(\boldsymbol{U}_0, \boldsymbol{U}_\star)$ given $\boldsymbol{Y}_\star$ is

$$\begin{pmatrix} \boldsymbol{\Sigma}_0 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{\Sigma_F} \end{pmatrix} - \begin{pmatrix} \boldsymbol{\Sigma}_0 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{\Sigma_F} \end{pmatrix} \begin{pmatrix} \boldsymbol{V}_0^T \\ \boldsymbol{V}_\star^T \end{pmatrix} \boldsymbol{\Sigma}_\star^{-1}(\boldsymbol{V}_0, \boldsymbol{V}_\star) \begin{pmatrix} \boldsymbol{\Sigma}_0 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{\Sigma_F} \end{pmatrix}. \quad \text{(S.7)}$$

Using the Woodbury matrix identity, we can further simplify $(\boldsymbol{V}_0, \boldsymbol{V}_\star)^T\boldsymbol{\Sigma}_\star^{-1}(\boldsymbol{V}_0, \boldsymbol{V}_\star)$ as

$$\begin{pmatrix} \boldsymbol{V}_0^T \\ \boldsymbol{V}_\star^T \end{pmatrix} \boldsymbol{\Sigma}_\star^{-1}(\boldsymbol{V}_0, \boldsymbol{V}_\star) = \boldsymbol{\Delta} - \boldsymbol{\Delta}\left[\begin{pmatrix} \boldsymbol{\Sigma}_0^{-1} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{\Sigma_F}^{-1} \end{pmatrix} + \boldsymbol{\Delta}\right]^{-1}\boldsymbol{\Delta},$$

where $\boldsymbol{\Delta} = (\boldsymbol{V}_0, \boldsymbol{V}_\star)^T \boldsymbol{\Sigma}_{\boldsymbol{E}}^{-1} (\boldsymbol{V}_0, \boldsymbol{V}_\star)$. Namely, at most we need to invert a $\sum_{k=0}^{K} r_k \times \sum_{k=0}^{K} r_k$ matrix, which is computationally feasible even for high dimensional data. In particular, under the orthogonal conditions, $\boldsymbol{\Delta}$ is a diagonal matrix:

$$\boldsymbol{\Delta} = \text{blkdiag}\left(\sum_{k=1}^{K} (\sigma_k^{-2}/K) \boldsymbol{I}_{r_0}, \sigma_1^{-2} \boldsymbol{I}_{r_1}, \cdots, \sigma_K^{-2} \boldsymbol{I}_{r_K} \right).$$

The computation is further simplified. All in all, the conditional distribution of the latent variables is fully derived. It is straightforward to calculate any conditional moment functions of $(\boldsymbol{U}_0, \boldsymbol{U}_\star)$. Therefore, we are ready for the M step.

B.2 M Step (Under the General Conditions)

In the M step, we maximize the conditional expectation of the joint log likelihood of the observed data and the latent variables. It can be separated into two sets of optimization problems: maximizing the conditional expectation of the marginal likelihood of $\boldsymbol{U}_k$'s $(k = 0, \cdots, K)$, and maximizing the conditional expectation of the conditional likelihood of $\boldsymbol{Y}_\star$ given $(\boldsymbol{U}_0, \boldsymbol{U}_\star)$.

First Optimization: The first set of optimization is relatively straightforward. Essentially, we need to solve the following problem

$$\min_{\boldsymbol{f}_k(\cdot), \boldsymbol{\Sigma}_k} \; n \log |\boldsymbol{\Sigma}_k| + \mathbb{E}_{\boldsymbol{U}_k | \boldsymbol{Y}_\star} \left\{ \text{tr} \left[(\boldsymbol{U}_k - \boldsymbol{f}_k(\boldsymbol{X})) \boldsymbol{\Sigma}_k^{-1} (\boldsymbol{U}_k - \boldsymbol{f}_k(\boldsymbol{X}))^T \right] \right\},$$

for $k = 0, 1, \cdots, K$. Since we assume $\boldsymbol{\Sigma}_k$ is diagonal and $\boldsymbol{f}_k = (f_{k,1}, \cdots, f_{k,r_k})$ contains r_k separate functions, the above problem can be further separated as

$$\min_{f_{k,j}(\cdot)} \; \mathbb{E}_{\boldsymbol{U}_k | \boldsymbol{Y}_\star} \left\{ \text{tr} \left[(\boldsymbol{u}_{k,j} - f_{k,j}(\boldsymbol{X}))(\boldsymbol{u}_{k,j} - f_{k,j}(\boldsymbol{X}))^T \right] \right\}, \quad j = 1, \cdots, r_k; \tag{S.8}$$

and

$$\min_{\boldsymbol{\Sigma}_k} \; n \log |\boldsymbol{\Sigma}_k| + \mathbb{E}_{\boldsymbol{U}_k | \boldsymbol{Y}_\star} \left\{ \text{tr} \left[(\boldsymbol{U}_k - \boldsymbol{f}_k(\boldsymbol{X})) \boldsymbol{\Sigma}_k^{-1} (\boldsymbol{U}_k - \boldsymbol{f}_k(\boldsymbol{X}))^T \right] \right\}. \tag{S.9}$$

After adding and subtracting some constant terms, (S.8) is equivalent to the least square problem:

$$\min_{f_{k,j}(\cdot)} \; \| \mathbb{E}(\boldsymbol{u}_{k,j} | \boldsymbol{Y}_\star) - f_{k,j}(\boldsymbol{X}) \|_{\mathbb{F}}^2, \tag{S.10}$$

for $j = 1, \cdots, r_k$. The optimization problem (S.10) has been well studied in the literature, and can be readily solved via off-the-shelf solvers. More specifically, if $f_{k,j}$ is a linear function, it becomes ordinary least squares

$$\min_{\boldsymbol{\beta}_{k,j}} \|\mathbb{E}(\boldsymbol{u}_{k,j}|\boldsymbol{Y}_\star) - \boldsymbol{X}\boldsymbol{\beta}_{k,j}\|_{\mathbb{F}}^2, \tag{S.11}$$

which has an explicit solution $\widehat{\boldsymbol{\beta}_{k,j}} = (\boldsymbol{X}^T\boldsymbol{X})^{-1}\boldsymbol{X}\mathbb{E}(\boldsymbol{u}_{k,j}|\boldsymbol{Y}_\star)$. If $f_{k,j}$ is nonparametric, (S.10) can be solved using kernel methods or spline-based methods (Hastie and Tibshirani, 1990; Fan and Gijbels, 1996; Hollander et al., 2013).

It is also straightforward to incorporate variable selection in (S.10) if $f_{k,j}$ is additive. For example, if $f_{k,j}$ is linear as in (S.11), we can add a sparsity-inducing penalty, such as the LASSO penalty (Tibshirani, 1996; Efron et al., 2004), to the objective function (S.11), and solve the penalized least squares problem to get a sparse estimate of $\boldsymbol{\beta}_{k,j}$. Tuning parameters can be selected adaptively. We particularly apply this procedure to the Berkeley Growth Study example in Section F. The results are more interpretable after variable selection. More generally, if $f_{k,j}$ is some nonparametric additive function, one may replace (S.10) with the methods for estimating sparse additive models proposed in the literature (cf. Lafferty and Wasserman, 2008; Ravikumar et al., 2009). All in all, incorporating variable selection into (S.10) is a well studied problem. With variable selection, we can accommodate high dimensional covariates and identify important covariates for each latent factor.

Once (S.10) is solved, the optimization problem (S.9) can be solved explicitly as

$$\widehat{\boldsymbol{\Sigma}_k} = \frac{1}{n}\mathrm{diag}\left\{\mathbb{E}_{\boldsymbol{U}_k|\boldsymbol{Y}_\star}\left[\left(\boldsymbol{U}_k - \widehat{\boldsymbol{f}_k}(\boldsymbol{X})\right)^T\left(\boldsymbol{U}_k - \widehat{\boldsymbol{f}_k}(\boldsymbol{X})\right)\right]\right\},\ k = 0, 1, \cdots, K \tag{S.12}$$

where $\widehat{\boldsymbol{f}_k} = (\widehat{f_{k,1}}, \cdots, \widehat{f_{k,r_k}})$ is the optimizer of (S.10).

Second Optimization: The second set of optimization can be written as

$$\min_{\boldsymbol{V}_0, \boldsymbol{V}_\star, \sigma_1^2, \cdots, \sigma_K^2} \sum_{k=1}^{K}\left[np_k \log\sigma_k^2 + \sigma_k^{-2}\mathbb{E}_{\boldsymbol{U}_0, \boldsymbol{U}_k|\boldsymbol{Y}_\star}\|\boldsymbol{Y}_k - \boldsymbol{U}_k\boldsymbol{V}_k^T - \boldsymbol{U}_0\boldsymbol{V}_{0,k}^T\|_{\mathbb{F}}^2\right], \tag{S.13}$$

where $\boldsymbol{V}_0$ and $\boldsymbol{V}_\star$ satisfy the basic conditions and the general conditions. Under those conditions, there are no closed-form expressions for $\boldsymbol{V}_0$ and $\boldsymbol{V}_k$ ($k = 1, \cdots, K$). As a remedy, we solve the above problem with respect to $\boldsymbol{V}_k$ and $\boldsymbol{V}_0$ sequentially.

Fixing $\boldsymbol{V}_0$, the optimization of $\boldsymbol{V}_k$ becomes

$$\begin{aligned}
&\min_{\boldsymbol{V}_k:\boldsymbol{V}_k^T\boldsymbol{V}_k=\boldsymbol{I}} \ \mathbb{E}_{\boldsymbol{U}_0,\boldsymbol{U}_k|\boldsymbol{Y}_\star}\|\boldsymbol{Y}_k-\boldsymbol{U}_0\boldsymbol{V}_{0,k}^T-\boldsymbol{U}_k\boldsymbol{V}_k^T\|_{\mathbb{F}}^2\\
\Leftrightarrow\ &\min_{\boldsymbol{V}_k:\boldsymbol{V}_k^T\boldsymbol{V}_k=\boldsymbol{I}} \ \mathbb{E}_{\boldsymbol{U}_0,\boldsymbol{U}_k|\boldsymbol{Y}_\star}\left\{-2\mathrm{tr}\left[\left(\boldsymbol{Y}_k-\boldsymbol{U}_0\boldsymbol{V}_{0,k}^T\right)^T\boldsymbol{U}_k\boldsymbol{V}_k^T\right]\right\}\\
\Leftrightarrow\ &\max_{\boldsymbol{V}_k:\boldsymbol{V}_k^T\boldsymbol{V}_k=\boldsymbol{I}} \ \mathrm{tr}\left[\mathbb{E}\left(\boldsymbol{Y}_k^T\boldsymbol{U}_k-\boldsymbol{V}_{0,k}\boldsymbol{U}_0^T\boldsymbol{U}_k|\boldsymbol{Y}_\star\right)\boldsymbol{V}_k^T\right]\\
\Leftrightarrow\ &\max_{\boldsymbol{V}_k:\boldsymbol{V}_k^T\boldsymbol{V}_k=\boldsymbol{I}} \ \mathrm{tr}\left\{\left[\boldsymbol{Y}_k^T\mathbb{E}(\boldsymbol{U}_k|\boldsymbol{Y}_\star)-\boldsymbol{V}_{0,k}\mathbb{E}(\boldsymbol{U}_0^T\boldsymbol{U}_k|\boldsymbol{Y}_\star)\right]\boldsymbol{V}_k^T\right\}. \qquad \text{(S.14)}
\end{aligned}$$

By the singular value decomposition (SVD), we have $\boldsymbol{Y}_k^T\mathbb{E}(\boldsymbol{U}_k|\boldsymbol{Y}_\star)-\boldsymbol{V}_{0,k}\mathbb{E}(\boldsymbol{U}_0^T\boldsymbol{U}_k|\boldsymbol{Y}_\star)=\boldsymbol{L}\boldsymbol{D}\boldsymbol{R}^T$, where $\boldsymbol{L}$ ($\boldsymbol{R}$) contains r_k left (right) singular vectors, and $\boldsymbol{D}$ is a diagonal matrix with r_k positive singular values on the diagonal. The object function in (S.14) satisfies the following relations

$$\mathrm{tr}(\boldsymbol{L}\boldsymbol{D}\boldsymbol{R}^T\boldsymbol{V}_k^T)=\mathrm{tr}(\boldsymbol{D}\boldsymbol{R}^T\boldsymbol{V}_k^T\boldsymbol{L})\le\mathrm{tr}(\boldsymbol{D}),$$

where "=" holds for the second inequality if and only if $\boldsymbol{V}_k=\boldsymbol{L}\boldsymbol{R}^T$. This is because $\boldsymbol{V}_k\boldsymbol{R}$ and $\boldsymbol{L}$ both have r_k orthonormal columns. The diagonal values of their inner product are no larger than 1. Since $\boldsymbol{D}$ has r_k positive diagonal values, $\mathrm{tr}(\boldsymbol{D}\boldsymbol{R}^T\boldsymbol{V}_k^T\boldsymbol{L})$ is maximized if and only if $\boldsymbol{V}_k\boldsymbol{R}=\boldsymbol{L}$. Namely, the optimal solution for (S.14) is

$$\widehat{\boldsymbol{V}_k}=\boldsymbol{L}\boldsymbol{R}^T. \qquad \text{(S.15)}$$

When $\boldsymbol{V}_k$'s are held fixed, the optimization of $\boldsymbol{V}_0$ becomes

$$\min_{\boldsymbol{V}_0:\boldsymbol{V}_0^T\boldsymbol{V}_0=\boldsymbol{I}} \ \sum_{k=1}^{K}\sigma_k^{-2}\mathbb{E}_{\boldsymbol{U}_0,\boldsymbol{U}_k|\boldsymbol{Y}_\star}\|\boldsymbol{Y}_k-\boldsymbol{U}_k\boldsymbol{V}_k^T-\boldsymbol{U}_0\boldsymbol{V}_{0,k}^T\|_{\mathbb{F}}^2. \qquad \text{(S.16)}$$

There is no closed-form solution for this constrained optimization problem. As a remedy, we use a relax-and-retrieve strategy to approximately solve the problem: first relax the orthogonality constraint and derive a closed-form solution for $\boldsymbol{V}_0$, and then retrieve the orthogonality through the eigendecomposition of $\boldsymbol{V}_0\boldsymbol{\Sigma}_0\boldsymbol{V}_0^T$. Without the constraint, the solution of (S.16) is

$$\widetilde{\boldsymbol{V}_{0,k}}=\left[\boldsymbol{Y}_k^T\mathbb{E}(\boldsymbol{U}_0|\boldsymbol{Y}_\star)-\boldsymbol{V}_k\mathbb{E}(\boldsymbol{U}_k^T\boldsymbol{U}_0|\boldsymbol{Y}_\star)\right]\left[\mathbb{E}(\boldsymbol{U}_0^T\boldsymbol{U}_0|\boldsymbol{Y}_\star)\right]^{-1},\ k=1,\cdots,K. \qquad \text{(S.17)}$$

Subsequently, we orthogonalize the columns in $\widetilde{\boldsymbol{V}_0}$ through the eigendecomposition of $\widetilde{\boldsymbol{V}_0}\widehat{\boldsymbol{\Sigma}_0}\widetilde{\boldsymbol{V}_0}^T$. In particular, we set the columns of $\widehat{\boldsymbol{V}_0}$ to be the eigenvectors, and

replace the original diagonal values of $\widehat{\boldsymbol{\Sigma}_0}$ with the eigenvalues of $\widetilde{\boldsymbol{V}_0}\widehat{\boldsymbol{\Sigma}_0}\widetilde{\boldsymbol{V}_0}^T$. This step will not change the value of $\boldsymbol{\Sigma}_\star$ in (S.2). A similar approach has been adopted in Li et al. (2016).

Once $\widehat{\boldsymbol{V}_0}$ and $\widehat{\boldsymbol{V}_k}$ are obtained, it is straightforward to optimize (S.13) to obtain the estimator of σ_k^2 as

$$\widehat{\sigma_k^2} = \mathbb{E}_{\boldsymbol{U}_0,\boldsymbol{U}_k|\boldsymbol{Y}_\star}\|\boldsymbol{Y}_k - \boldsymbol{U}_k\widehat{\boldsymbol{V}_k}^T - \boldsymbol{U}_0\widehat{\boldsymbol{V}_{0,k}}^T\|_{\mathbb{F}}^2, \quad k = 1, \cdots, K. \tag{S.18}$$

B.3 M Step (Under the Orthogonal Conditions)

Under the orthogonal conditions, the first set of optimization remains the same with those under the general conditions. We only show a variant of the algorithm for the second set of optimization.

With the orthogonal constraint $(\sqrt{K}\boldsymbol{V}_{0,k}, \boldsymbol{V}_k)^T(\sqrt{K}\boldsymbol{V}_{0,k}, \boldsymbol{V}_k) = \boldsymbol{I}$ for $k = 1, \cdots, K$, the original optimization problem (S.13) can be separated into K parts and solved in parallel. In particular, $\boldsymbol{V}_{0,k}$ and $\boldsymbol{V}_k$ can be estimated together via solving the following optimization under the orthogonal constraint

$$\begin{aligned}
&\min_{\boldsymbol{V}_{0,k},\boldsymbol{V}_k} \ \mathbb{E}_{\boldsymbol{U}_0,\boldsymbol{U}_k|\boldsymbol{Y}_\star}\|\boldsymbol{Y}_k - \boldsymbol{U}_k\boldsymbol{V}_k^T - \boldsymbol{U}_0\boldsymbol{V}_{0,k}^T\|_{\mathbb{F}}^2 \\
\Leftrightarrow \ &\min_{\boldsymbol{V}_{0,k},\boldsymbol{V}_k} \ \mathbb{E}_{\boldsymbol{U}_0,\boldsymbol{U}_k|\boldsymbol{Y}_\star}\|\boldsymbol{Y}_k - (\frac{1}{\sqrt{K}}\boldsymbol{U}_0, \boldsymbol{U}_k)(\sqrt{K}\boldsymbol{V}_{0,k}, \boldsymbol{V}_k)^T\|_{\mathbb{F}}^2 \\
\Leftrightarrow \ &\min_{\boldsymbol{V}_{0,k},\boldsymbol{V}_k} \ \|\boldsymbol{Y}_k - \left(\frac{1}{\sqrt{K}}\mathbb{E}(\boldsymbol{U}_0|\boldsymbol{Y}_\star), \mathbb{E}(\boldsymbol{U}_k|\boldsymbol{Y}_\star)\right)(\sqrt{K}\boldsymbol{V}_{0,k}, \boldsymbol{V}_k)^T\|_{\mathbb{F}}^2.
\end{aligned} \tag{S.19}$$

The above problem is exactly in the form of an orthogonal Procrustes problem (Gower and Dijksterhuis, 2004). Thus the unique optimal solution for (S.19) is

$$(\sqrt{K}\widehat{\boldsymbol{V}_{0,k}}, \widehat{\boldsymbol{V}_k}) = \boldsymbol{L}\boldsymbol{R}^T, \tag{S.20}$$

where $\boldsymbol{L}$ and $\boldsymbol{R}$ contain the r_0+r_k left and right singular vectors of $\boldsymbol{Y}_k^T\left(1/\sqrt{K}\mathbb{E}(\boldsymbol{U}_0|\boldsymbol{Y}_\star), \mathbb{E}(\boldsymbol{U}_k|\boldsymbol{Y}_\star)\right)$, respectively. Subsequently, we obtain $\widehat{\boldsymbol{V}_0}$ by concatenating the respective $\widehat{\boldsymbol{V}_{0,k}}$'s. Once the loading matrices are estimated, we use the same estimator as before to estimate σ_k^2.

B.4 Flow Chart

We summarize the EM algorithms for model fitting under different identifiability conditions in Algorithm 1 and Algorithm 2. We emphasize that under the orthogonal

conditions, the computation is extremely efficient. There is no need for sequential optimization or approximation.

Algorithm 1 The EM algorithm under the *general conditions*

Initialize θ_0, possibly from JIVE or PCA estimates ;

while Estimation has not reached convergence **do**

E Step:

Calculate the conditional mean (S.6) and conditional variance (S.7)

M Step:

Estimate $f_{k,j}$ $(j = 1, \cdots, r_k; k = 0, 1, \cdots, K)$ by solving (S.10);

Estimate $\boldsymbol{\Sigma}_k$ $(k = 0, 1, \cdots, K)$ via (S.12);

Cycle through the following steps (for one round or until convergence):

- Estimate $\boldsymbol{V}_k$ $(k = 1, \cdots, K)$ while fixing $\boldsymbol{V}_0$ via (S.15);
- Obtain an intermediate estimate of $\boldsymbol{V}_{0,k}$ while fixing $\boldsymbol{V}_k$ $(k = 1, \cdots, K)$ via (S.17);
- Normalize the intermediate estimate of $\boldsymbol{V}_0$ and $\boldsymbol{\Sigma}_0$ together through SVD;

Estimate σ_k^2 $(k = 1, \cdots, K)$ via (S.18);

end while

C Likelihood Cross Validation for Refining Rank Estimation

In the SIFA model, a set of the joint rank r_0 and the individual ranks r_k $(k = 1, \cdots, K)$ is treated as a tuning parameter, denoted by $\boldsymbol{r} = (r_0, r_1, \cdots, r_K)$. Given a collection of candidate rank sets, we introduce an N-fold LCV method to select the best rank set. More specifically, we first randomly split the samples across all data sets into N groups. In each cross validation run, one block of samples is treated as testing samples, and the other samples are used as training samples. Suppose we have m candidate rank sets. In each run, we fit m models under different ranks using the training samples. Then we calculate the log likelihood value of the testing samples

Algorithm 2 The EM algorithm under the *orthogonal conditions*

Initialize θ_0, possibly from JIVE or PCA estimates ;

while Estimation has not reached convergence **do**

E Step:

Calculate the conditional mean (S.6) and conditional variance (S.7)

M Step:

Estimate $f_{k,j}$ $(j = 1, \cdots, r_k; k = 0, 1, \cdots, K)$ by solving (S.10);

Estimate $\mathbf{\Sigma}_k$ $(k = 0, 1, \cdots, K)$ via (S.12);

Estimate $\boldsymbol{V}_{0,k}$ and $\boldsymbol{V}_k$ $(k = 1, \cdots, K)$ via (S.20)

Estimate σ_k^2 $(k = 1, \cdots, K)$ via (S.18);

end while

for each fitted model, using Eqn (7) in the main paper. The negative log likelihood values are used as LCV scores, where a smaller value is more desired. We repeat the procedure N times and average the LCV scores for each candidate rank set. The rank set corresponds to the smallest score is selected. We remark that an N-fold LCV study on M candidate rank sets requires fitting the SIFA model $M * N$ times. This may be computationally intensive when N or M is large. Therefore, in practice, when the data dimension is large, we recommend estimating the ranks with the two-step procedure first, and then using the LCV approach to refine the result by searching among the neighbors of the crude estimate.

To demonstrate the efficacy of the LCV approach, we consider a couple of simulated examples. In particular, data are generated from Setting 1 and 2, respectively, in Section 3.1 of the main paper, with the true ranks to be $r_0 = 2, r_1 = r_2 = 3$. We consider 9 rank sets: $(r_0, r_1, r_2) \in \{(1,2,2), (2,2,2), (3,2,2), (1,3,3), (2,3,3), (3,3,3), (3,4,3), (3,4,4), (4,4,4)\}$, in the neighborhood of the true rank set. The 10-fold LCV scores (i.e., negative log likelihood values) are shown in Figure S5 (for Setting 2) and Figure S6 (for Setting 1), respectively.

When the data are generated from Setting 2 (i.e., the SIFA-A model), the LCV approach correctly selects the true rank set. We remark that the sets with large ranks all have relatively small LCV scores, so there may be the risk of overestimation. When the data are generated from Setting 1 (i.e., the JIVE model), the LCV approach does

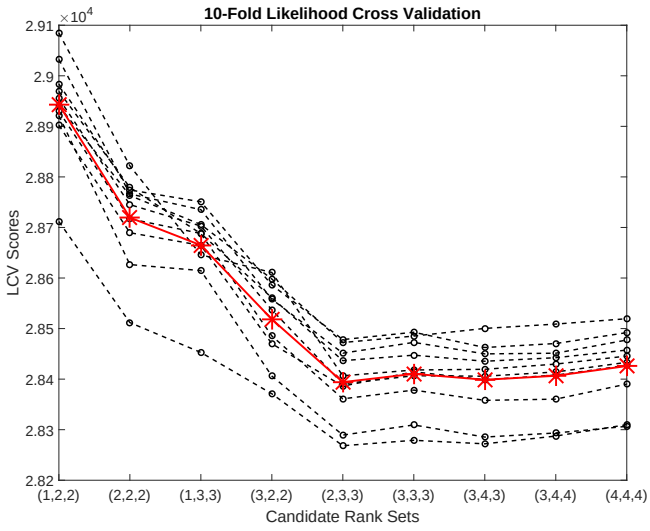

Figure S5: Setting 2: The LCV scores for 10-fold cross validation on 9 candidate rank sets. Each dashed line with circles contains corresponds to the LCV scores (negative log likelihood values) in one fold. The solid line with stars contains the average LCV scores for different rank sets.

not correctly select the ranks. The CV scores for different rank sets are quite similar. This may be due to the lack of relevant information from auxiliary covariates. A more powerful and customized rank estimation procedure requires further investigation.

D Additional Simulation Studies

D.1 Nonlinear Covariate Functions

To investigate nonparametric model fitting under nonlinear covariate functions, we consider the following simulation setting. In particular, the parameters and the ranks are set in the same way as in Setting 2 in the main manuscript, except there is only one covariate ($q = 1$) and it is related to different latent factors in highly nonlinear fashions.

- **Setting 4** (SIFA-A Model with nonlinear relations): The factors are generated from $\boldsymbol{U}_k = \boldsymbol{f}_k(X) + \boldsymbol{F}_k$ for $k = 0, 1, 2$, where X represents the univariate covariate. The univariate functions in $\boldsymbol{f}_0(\cdot), \boldsymbol{f}_1(\cdot), \boldsymbol{f}_2(\cdot)$ contain sine, cosine,

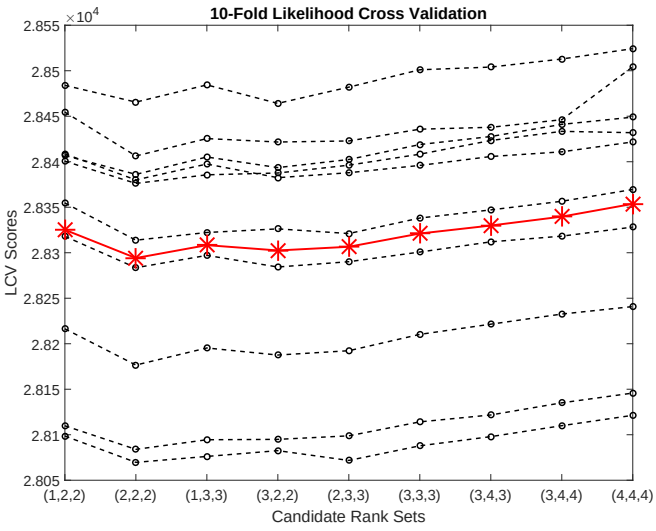

Figure S6: Setting 1: The LCV scores for 10-fold cross validation on 9 candidate rank sets. Each dashed line with circles contains corresponds to the LCV scores (negative log likelihood values) in one fold. The solid line with stars contains the average LCV scores for different rank sets.

quadratic, and cubic functions. Other parameters are generated in compliance with the general conditions.

We focus on SIFA-A models with nonparametric estimation and linear estimation. The nonparametric estimation is achieved using a kernel regression method. In addition, we also consider the JIVE model without incorporating the covariate. The comparison results are shown in Figure S7. In terms of the estimation accuracy for separate loadings and the combined loading subspace, it is apparent that using the kernel regression method in SIFA-A achieves the best result. The SIFA-A model with misspecified linear covariate functions has comparable performance with JIVE. For the low-rank structure recovery accuracy, both SIFA-A models significantly outperform the JIVE model, with the nonparametric version being slightly better than the linear version. Overall, incorporating the auxiliary covariate into the model improves parameter estimation and dimension reduction. SIFA is also robust against covariate function misspecification.

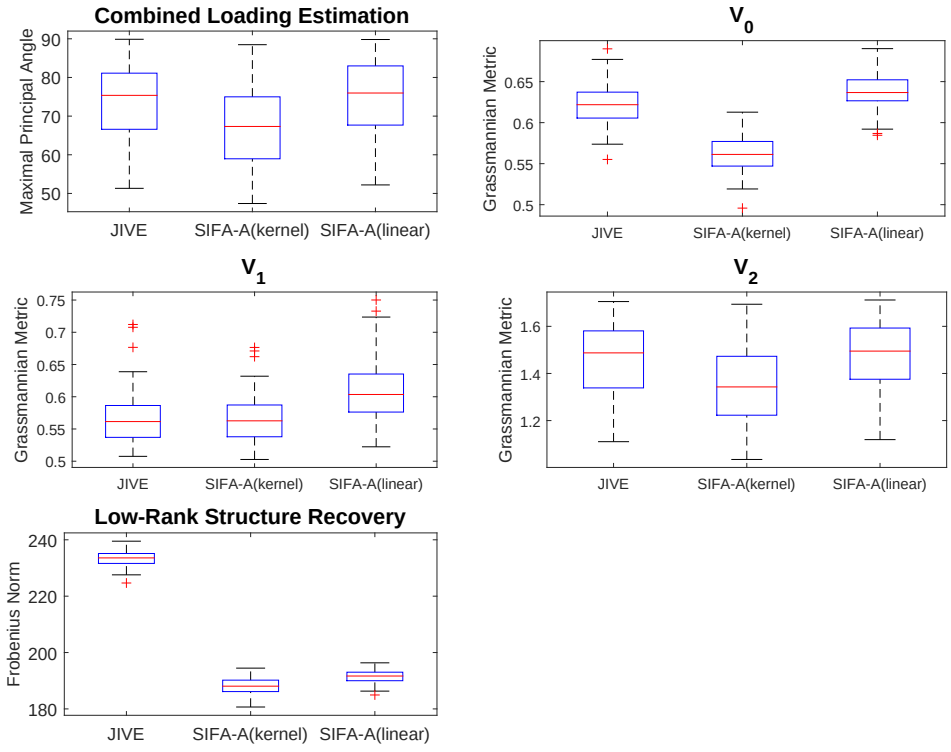

Figure S7: Simulation study of nonlinear covariate functions. Upper left: the maximal principal angles between the true and estimated loadings under Setting 4; Upper right and middle two: the Grassmannian metric of the differences between the true and estimated loading matrices; Lower left: the Frobenius norm of the differences between the true and estimated low-rank structure. In each figure, from left to right, the box plots correspond to JIVE, SIFA-A with kernel regression, and SIFA-A with linear regression, respectively. The results are based on 100 simulation runs.

D.2 Overfitting of Nonparametric Estimation

We also study the overfitting of nonparametric estimation in SIFA models when the true covariate functions are linear. In particular, we consider the following simulation setting where the parameters and the ranks are set in the same way as in Setting 3 in the main manuscript, except there is only one covariate ($q = 1$).

- **Setting 5** (SIFA-B Model with univariate linear relations): The latent factors are generated from $\boldsymbol{U}_k = X\boldsymbol{b}_k^T + \boldsymbol{F}_k$ for $k = 0, 1, 2$, where X represents a univariate covariate and $\boldsymbol{b}_k$ is a length-r_k coefficient vector. Other parameters are generated in compliance with the orthogonal conditions.

We focus on SIFA-B methods with nonparametric and linear relations, and the JIVE method. The comparison results are shown in Figure S8. SIFA-B fitted with linear regressions (i.e., the true model) has the best estimation accuracy, both for loadings and the low-rank structure. SIFA-B fitted with nonparametric regressions has very similar but slightly worse performance, probably due to the overfitting issue. Both methods significantly outperform the JIVE method. It indicates the overfitting issue of the nonparametric covariate functions is marginal compared to the exclusion of the auxiliary information (as in JIVE).

D.3 Rank Misspecification

In this section, we provide simulation results for the rank misspecification study. Data are generated from the Setting 2 (i.e., SIFA-A model). The true ranks are $r_0 = 2, r_1 = r_2 = 3$.

First, we overestimate every rank by 1. Namely, misspecified ranks $\widehat{r_0} = 3, \widehat{r_1} = \widehat{r_2} = 4$ are used for fitting JIVE, SIFA-A and SIFA-B, and $\widehat{r_0} + \widehat{r_1} + \widehat{r_2} = 11$ is used as the single rank for fitting PCA and SupSVD. The results are shown in Figure S9 and Figure S10. Similar to the simulation results in the main article for Setting 2, SIFA-A outperforms JIVE and SIFA-B in terms of the estimation accuracy of $\boldsymbol{V}_0$ and $\boldsymbol{V}_2$. The $\boldsymbol{V}_1$ estimate of SIFA-A has a relatively larger variance possibly because one of the individual patterns is captured by the redundant joint rank. Nevertheless, in the top panel of Figure S10, we can see that SIFA-A provides the best overall estimation

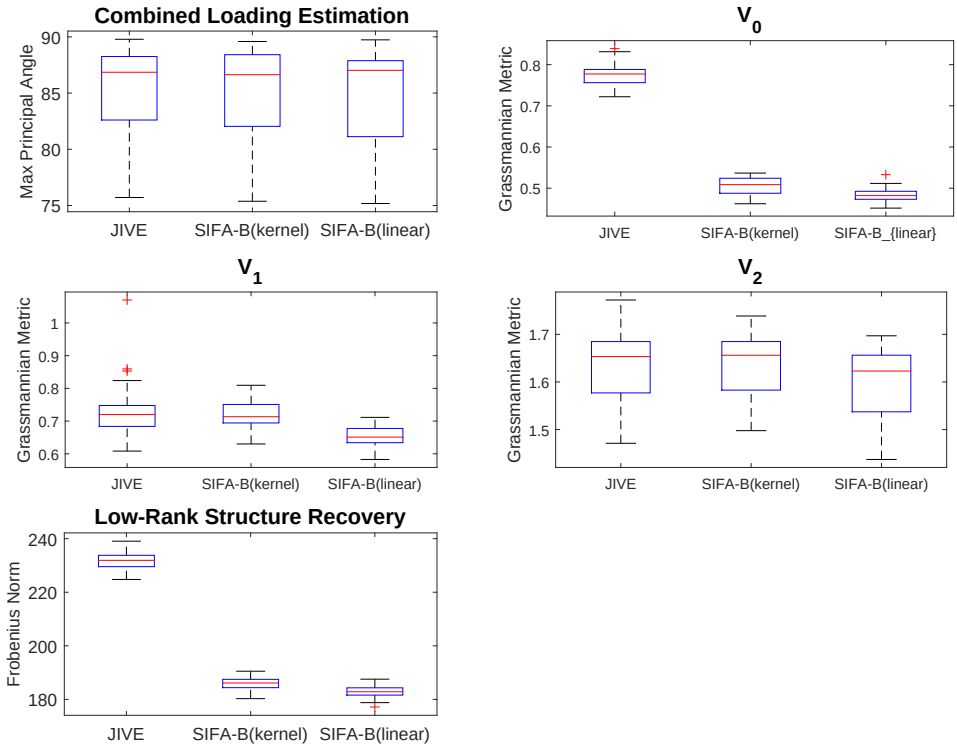

Figure S8: Simulation study of overfitting of nonparametric estimation. Upper left: the maximal principal angles between the true and estimated loadings under Setting 4; Upper right and middle two: the Grassmannian metric of the differences between the true and estimated loading matrices; Lower left: the Frobenius norm of the differences between the true and estimated low-rank structure. In each figure, from left to right, the box plots correspond to JIVE, SIFA-B with kernel regression, and SIFA-B with linear regression, respectively. The results are based on 100 simulation runs.

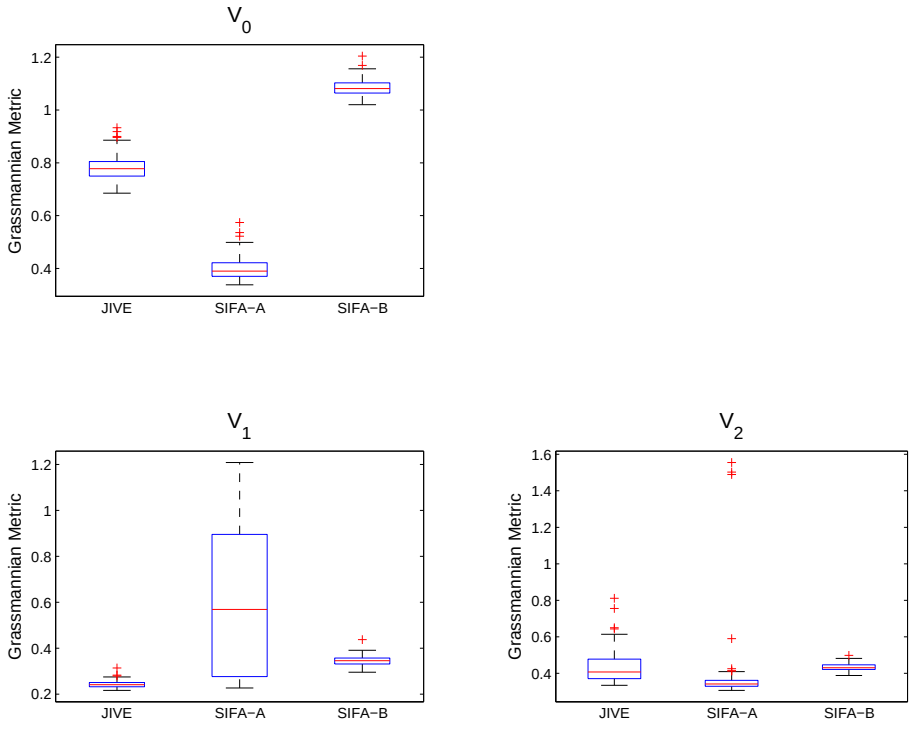

Figure S9: Simulation study of rank overestimation. The Grassmannian metric of the differences between true and estimated loading matrices under Setting 2 (the estimated loading matrices have overestimated ranks). In each figure, from left to right, the box plots correspond to JIVE, SIFA-A, and SIFA-B respectively. The results are based on 100 simulation runs.

accuracy for the combined loadings. It also has the best low-rank structure recovery accuracy among all methods.

Then we investigate the effect of underestimated ranks by setting $\widehat{r_0} = 1, \widehat{r_1} = \widehat{r_2} = 2$. The simulation results are shown in Figure S11 and Figure S12. Again, the results are similar to those obtained from the simulation study under Setting 2 in the main article. SIFA-A uniformly outperforms JIVE and SIFA-B in separate loading estimation, combined loading estimation and low-rank structure recovery. Interestingly, we note that in Figure S12, the more sophisticated integrative decomposition methods (i.e., JIVE, SIFA-A and SIFA-B) are worse than the simpler decomposition methods (i.e., PCA and SupSVD for concatenated data) with regard to the combined

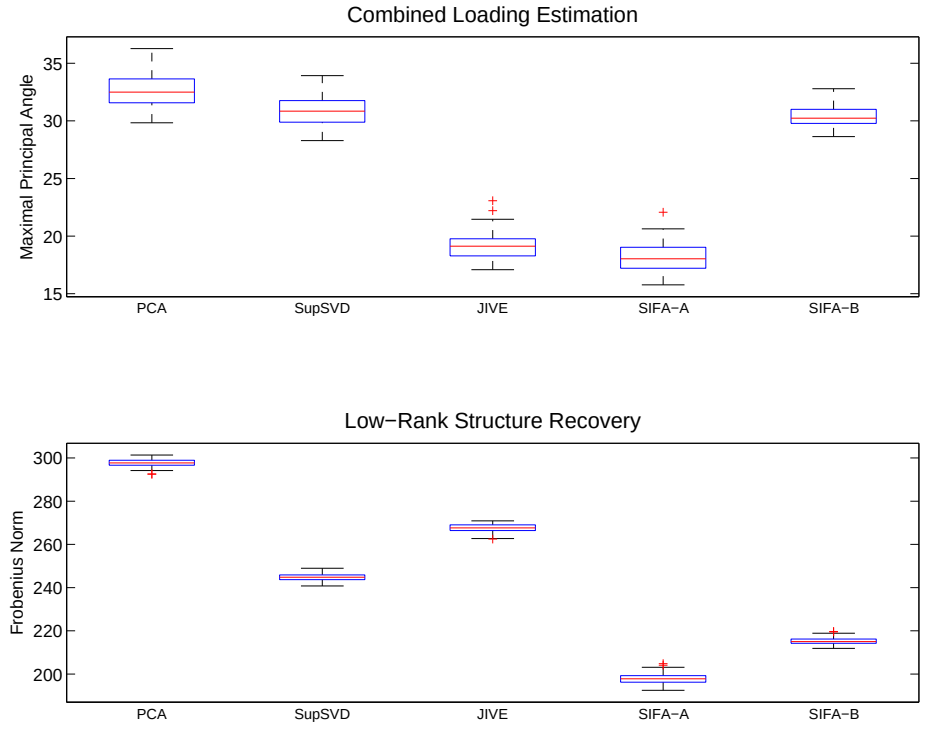

Figure S10: Simulation study of rank overestimation. Upper: the maximal principal angles between true and estimated loadings under Setting 2 (the estimated loading matrices have overestimated ranks); Lower: the Frobenius norm of the differences between true and estimated low-rank structures. In each figure, from left to right, the box plots correspond to PCA, SupSVD, JIVE, SIFA-A, and SIFA-B respectively. The results are based on 100 simulation runs.

loading estimation and low-rank structure recovery. This is mainly because the simpler decomposition methods do not impose any rank restriction for different patterns. They have the freedom to allocate ranks optimally to capture the underlying structure with the largest variation. When all the ranks are underestimated, the higher flexibility of rank allocation leads to the better overall performance. Nonetheless, the integrative decomposition methods (especially the proposed SIFA methods) have the unique capacity of distinguishing joint and individual patterns, which potentially enhance interpretation and facilitate further statistical analyses. Therefore, they are still deemed very useful in practice even when the ranks are misspecified.

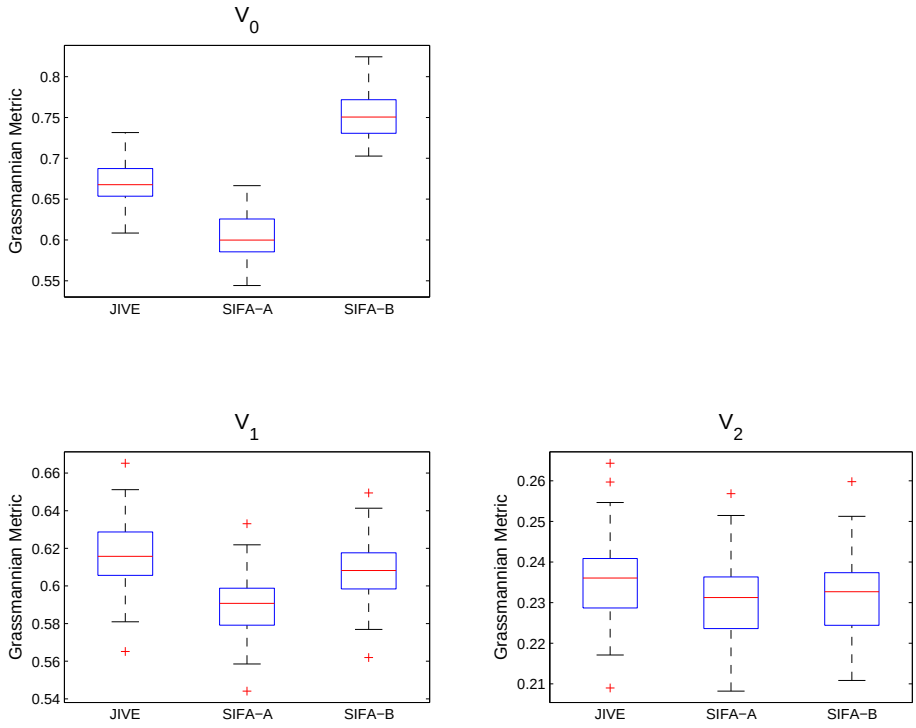

Figure S11: Simulation study of rank underestimation. The Grassmannian metric of the differences between true and estimated loading matrices under Setting 2 (the estimated loading matrices have underestimated ranks). In each figure, from left to right, the box plots correspond to JIVE, SIFA-A, and SIFA-B respectively. The results are based on 100 simulation runs.

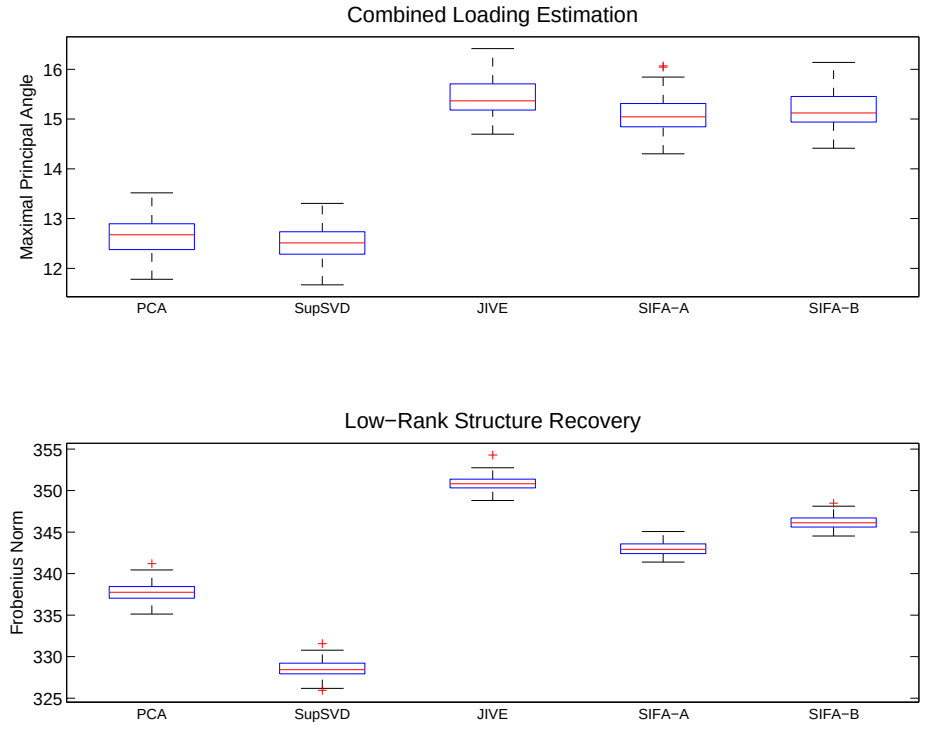

Figure S12: Simulation study of rank underestimation. Upper: the maximal principal angles between true and estimated loadings under Setting 2 (the estimated loading matrices have underestimated ranks); Lower: the Frobenius norm of the differences between true and estimated low-rank structures. In each figure, from left to right, the box plots correspond to PCA, SupSVD, JIVE, SIFA-A, and SIFA-B respectively. The results are based on 100 simulation runs.

D.4 Non-Gaussian Noise

In this section, we investigate the influence of the non-Gaussian noise on model fitting. We exploit the simulation Setting 3 in the main manuscript, except for the measurement error matrices $\boldsymbol{E}_1$ and $\boldsymbol{E}_2$. We assume the noise variables in $\boldsymbol{E}_1$ and $\boldsymbol{E}_2$ are i.i.d. from the same distribution, and consider 5 different distributions: the standard Gaussian distribution, and the t distributions with the degree of freedom 3, 5, 10, and 20, respectively. In particular, for each t distribution, we multiply the error matrices with a constant (i.e., $\sqrt{(df-2)/df}$) so that the random variables have unit variance. Then we fit the SIFA-B model, and evaluate the loading estimation and low-rank structure recovery accuracy under different settings. The results are presented in Figure S13. We remark that the estimation is not too much affected by the violation of the Gaussian assumption. When the noise distribution is very heavy-tailed (i.e., $t(3)/\sqrt{3}$), the variation of the estimation over different simulation runs increases. This is possibly due to the effect of the outliers in the noise matrix. All in all, the method is very robust when the noise is non-Gaussian.

D.5 Rescaling of Different Data Sets

Statistical decisions or actions based on data should not be affected by simple transformation. Here we first show that the SIFA models under the general conditions are equivariant under individual scaling of data sets, and extend the results to the case of highly-imbalanced dimensions of data sets. Suppose we have a sample of size n for $K = 2$ data sets from a SIFA model satisfying the orthogonal conditions:

$$\boldsymbol{Y}_1 = (\boldsymbol{f}_0(\boldsymbol{X}) + \boldsymbol{F}_0)\boldsymbol{V}_{0,1}^T + (\boldsymbol{f}_1(\boldsymbol{X}) + \boldsymbol{F}_1)\boldsymbol{V}_1^T + \boldsymbol{E}_1, \tag{S.21}$$

$$\boldsymbol{Y}_2 = (\boldsymbol{f}_0(\boldsymbol{X}) + \boldsymbol{F}_0)\boldsymbol{V}_{0,2}^T + (\boldsymbol{f}_2(\boldsymbol{X}) + \boldsymbol{F}_2)\boldsymbol{V}_2^T + \boldsymbol{E}_2. \tag{S.22}$$

If the measurement unit of $\boldsymbol{Y}_1$ has changed so that all entries of $\boldsymbol{Y}_1$ is scaled by some $s > 0$, then the combined model for the scaled data becomes

$$[s\boldsymbol{Y}_1; \boldsymbol{Y}_2] = c_s(\boldsymbol{f}_0(\boldsymbol{X}) + \boldsymbol{F}_0)\boldsymbol{V}_{0(s)}^T + [s\boldsymbol{f}_1(\boldsymbol{X}) + s\boldsymbol{F}_1; \boldsymbol{f}_2(\boldsymbol{X}) + \boldsymbol{F}_2]\text{blkdiag}(\boldsymbol{V}_1^T, \boldsymbol{V}_2^T) + [s\boldsymbol{E}_1; \boldsymbol{E}_2] \tag{S.23}$$

where $c_s = s^2/2 + 1/2$ and $\boldsymbol{V}_{0(s)}^T = [s\boldsymbol{V}_{0,1}^T; \boldsymbol{V}_{0,2}^T]/c_s$. This *scaled* model does not satisfy the orthogonal conditions any more, but it still satisfies the general conditions. Since

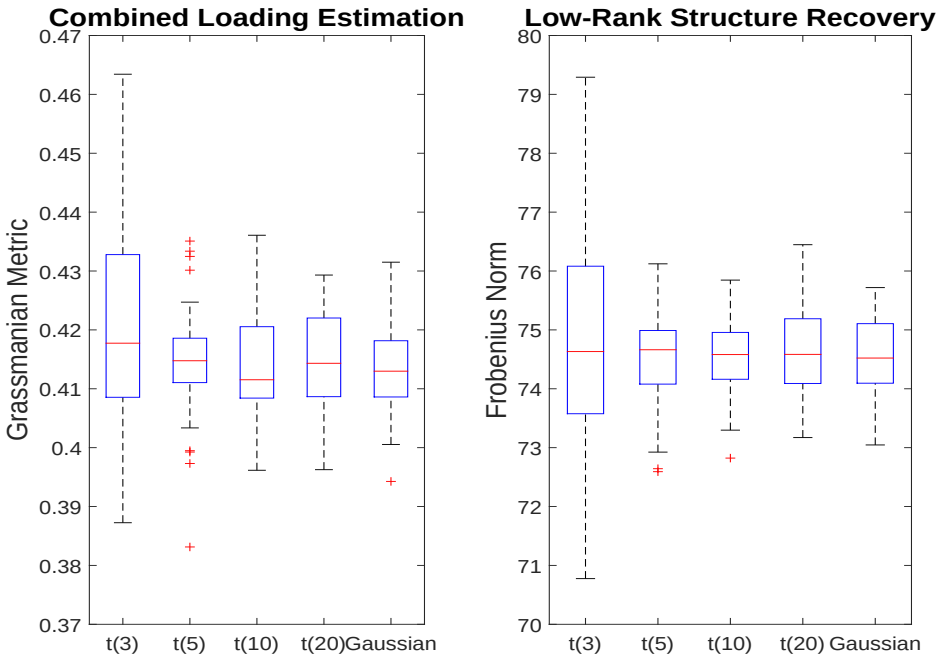

Figure S13: Simulation study of non-Gaussian noise. Left: the Grassmanian metric between the true and the estimated loadings; Right: the Frobenius norm of the differences between the true and the estimated low-rank structures. In each panel, from left to right, the box plots correspond to the noise distributions $t(3)/\sqrt{3},\ t(5)/\sqrt{5/3},\ t(10)/\sqrt{5/4},\ t(20)/\sqrt{10/9}$ and $\mathcal{N}(0,1)$.

for any value of s, the model consists of the same direction vectors $\boldsymbol{V}_{0,1}, \boldsymbol{V}_{0,2}, \boldsymbol{V}_1$ and $\boldsymbol{V}_2$ and the interpretation of the scores remains the same for any s, the SIFA model under the general conditions is equivariant under arbitrary scaling of individual data sets.

While the SIFA model under the general conditions is equivariant to individual scaling, the likelihood-based estimators can be sensitive to relative sizes of data sets. Nevertheless, we have observed that the SIFA estimates under general conditions, SIFA-A estimates, are *nearly equivariant*. We demonstrate the near-equivariance of SIFA-A estimates for a simulated data set, with a comparison to SIFA-B estimates and JIVE estimates. For this experiment, we have generated a sample from the SIFA model (S.21)-(S.22) satisfying the orthogonal conditions. For simplicity, we used rank 1 for each of joint and individual variations. The first data set $\boldsymbol{Y}_1$ is then individually rescaled by s ranging from 0.01 to 100; that is, in one extreme, the Frobenius norm of $\boldsymbol{Y}_1$ is approximately 100 times larger than that of $\boldsymbol{Y}_2$ (and vice versa). For this example, we check that the loading estimates of SIFA-A are nearly equivariant under individual scaling, as shown in Figures S14 and S15. On the other hand, the estimated loadings of SIFA-B and JIVE are quite dependent to the scale, and the estimation accuracy of those is worse than that of SIFA-A. Note that for any $s \neq 0$, the true model violates the orthogonal conditions of SIFA-B. The low-rank structure recovery of SIFA-A is also among the best across various scales (see Figure S14).

We further point out that the predicted scores from each rescaled data sets are highly correlated to each other. Moreover, as shown in Figure S15, the predictions of SIFA-A are highly correlated with the true scores with correlation coefficients > 0.8 for all values of s considered. Thus, any subsequent analysis using the score predictions of SIFA-A are invariant to individual scaling of data sets. On the other hand, JIVE does not possess such equivariance property and the correlation coefficients to their predictions are as low as 0.2. In this and many other examples, we have observed that SIFA-B score predictions were as good as those of SIFA-A. Although SIFA-B loading estimates are not as accurate as those of SIFA-A, they are at most about 50 degrees away from the truth. Heuristically, the projections of data points along the true $\mathbf{v}$ to $\hat{\mathbf{v}}$ preserve the order and relative magnitudes of original scores, if they are not orthogonal, which may explain the high score prediction accuracy of SIFA-

B. Overall, in terms of score prediction, SIFA-A and SIFA-B both have shown the invariance to individual scaling.

We also investigate the SIFA model fitting when the the dimensions of the individual data sets are highly imbalanced. We remark that the SIFA-B model is robust against imbalanced dimensions, because of the orthogonal and equal norm identifiability conditions. For the SIFA-A model, it is challenging to derive an argument similar to the equivariance under the rescaling of data. However, we observed the near-equivariance under two scenarios of increasing dimensions (of the first data set), while the second data set is kept fixed. The first scenario corresponds to the case where variables in the increasing dimension has factors highly correlated to existing factors, while the second scenario is the case where the added variables are pure noises. In each of both cases, we artificially added variables to make dimensions imbalanced (up to $(p_1, p_2) = (5000, 100)$). Notably, the proportion of $\|\hat{V}_{0,1}\|^2$ to $\|\hat{V}_{0,1}\|^2 + \|\hat{V}_{0,2}\|^2$ increases as p_1 increases (but at a very slow rate). In the meantime, the loading estimates of the SIFA-A model remain to be significantly better than the JIVE estimates (graphical results are omitted, but similar to Figures S10-S11) as p_1 increases. We remark that the first scenario of increasing dimensions produces a similar log-likelihood function as that from scaling the first data set by $s = p_1/p_2$. Thus, the near-equivariance is indeed expected.

D.6 Scalability of SIFA Algorithms

We also investigate the scalability of the proposed Algorithms 1 and 2 for model fitting. We exploit the simulation Setting 3 in the main manuscript and consider a range of dimensions to generate data. In particular, we consider 5 settings with the sample size n, the dimensions (p_1, p_2), and the number of covariates q being: $(n, p_1, p_2, q) = (100, 100, 100, 10), (200, 200, 200, 20), (500, 500, 500, 50),$
$(1000, 1000, 1000, 100), (2000, 2000, 2000, 200)$. For each setting, we conduct 30 simulation runs. In each simulation run, we fit SIFA-A and SIFA-B to the data. The computing time on a standard desktop is summarized in Figure S16. We remark that both algorithms are highly efficient. Even for thousands of samples and dimensions, it only takes a few minutes to fit a SIFA model on a desktop. The SIFA-B algorithm is more computationally efficient than the SIFA-A algorithm since the M step has

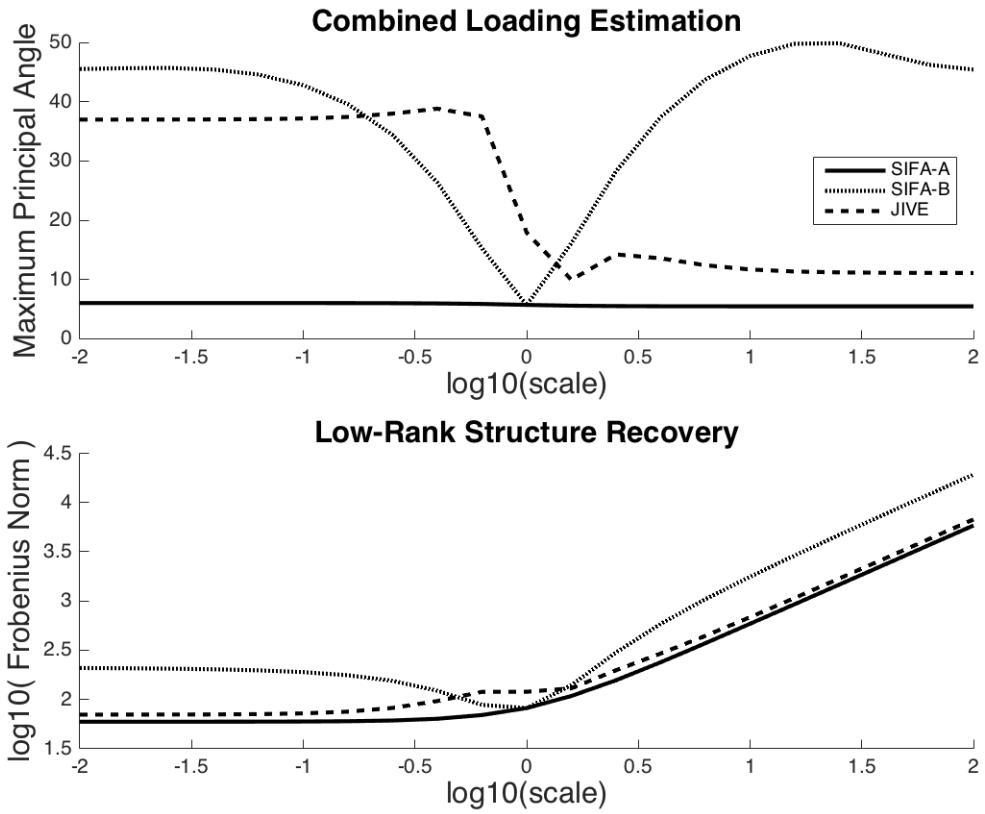

Figure S14: Simulation study for data scaling. Top: the maximal principal angles between the true ($[\boldsymbol{V}_{0(s)}, \text{blkdiag}(\boldsymbol{V}_1^T, \boldsymbol{V}_2^T)]$) and estimated loadings; Bottom: the Frobenius norm of the differences between the true and estimated low-rank structure. The true loading and low-rank structure for each scale value, s, are from (S.23).

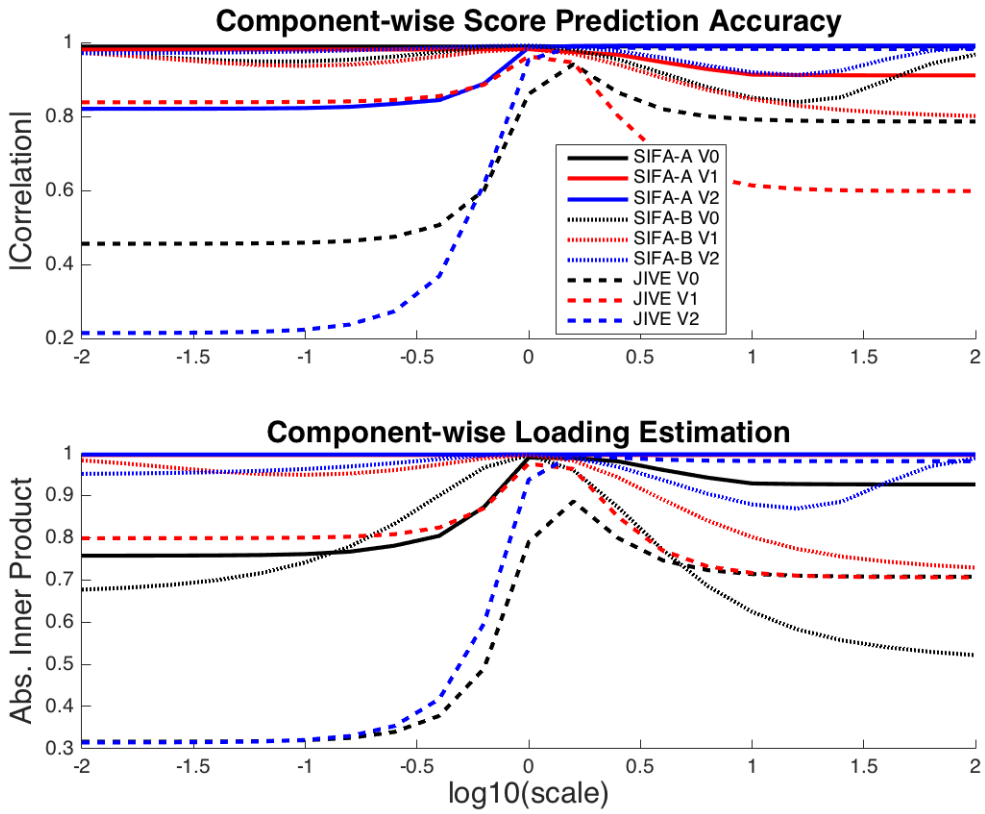

Figure S15: Simulation study for data scaling. Top: the absolute correlation coefficients between the true $\boldsymbol{f}_i(\boldsymbol{X}) + \boldsymbol{F}_i$ and its predictions ($i = 0, 1, 2$); Bottom: The inner product between the true and estimated loadings. The true loading and low-rank structure for each scale value, s, are from (S.23).

explicit solutions. The directions for further improvement include parallel computing and distributed computing.

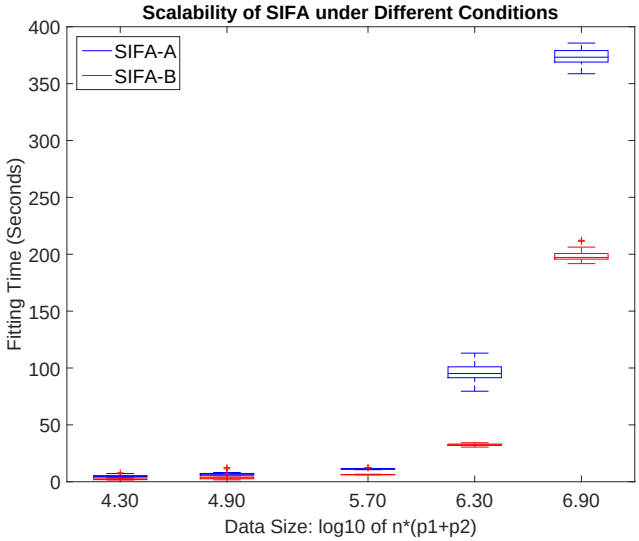

Figure S16: Scalability of the SIFA model fitting algorithms. From left to right, we consider simulation settings with the sample size-dimension tuples $(n, p_1, p_2, q) \in \{(100, 100, 100, 10),$
$(200, 200, 200, 20), (500, 500, 500, 50), (1000, 1000, 1000, 100), (2000, 2000, 2000, 200)\}$. The blue boxes correspond to the fitting times of SIFA-A while the red boxes correspond to the fitting times of SIFA-B. The x-axis is on the $\log_{10}$ scale of the total number of entries in the primary data sets. The result is based on 30 simulation runs in each setting.

E GTEx Data Analysis

In this section, we provide additional details on the data preprocessing and rank estimation of the GTEx example.

E.1 Data Preprocessing

We focus on three tissues with the most samples in the GTEx data, i.e., muscle ($n = 361$), blood ($n = 338$), and skin ($n = 302$). There are 204 common samples across 3 tissues. The p53 signaling pathway contains around 200 genes (the

gene names can be found on the GSEA website `http://software.broadinstitute.org/gsea/msigdb/cards/HALLMARK_P53_PATHWAY.html`). After removing the non-expressed genes according to the GTEx preprocessing criteria, we end up with 191 genes in each tissue. We normalize the expression level of each gene in each tissue through an inverse normal transformation as in Ardlie et al. (2015). Consequently, each gene follows a normal distribution with mean zero and unit variance. We denote the three preprocessed expression matrices as $\boldsymbol{Y}_1, \boldsymbol{Y}_2, \boldsymbol{Y}_3$, which are the primary input data for the SIFA method.

For each of the 204 common samples, we also have the sex information, the genotyping array platform index, and the individual's genotype information. The genotype data contain the mutation status of 6,856,774 single nucleotide polymorphisms (SNPs). We particularly restrict our scope to the *cis* SNPs of the genes in the p53 pathway, i.e., SNPs that lie within 1 megabase pair (1Mbp) of the transcription start site of a gene. Moreover, we eliminate those SNPs with minor allele frequency smaller than 10% across samples. As a result, we obtain 639,965 SNPs. The dimension of the genotype data far exceeds the sample size. To further reduce the dimension of the genotype data, we apply PCA to the data and use the top 30 PC score vectors to represent the genotype data. We note that the first PC score vector typically captures the ancestry information (see Figure S17) and the other PC scores capture the majority of variation in the SNP data. We column-center the two binary variables and the 30 PC score vectors and treat them as the covariates, denoted by $\boldsymbol{X}$, for the SIFA method.

Because our primary goal is to identify gene expression patterns across tissues and quantify the regulatory effect of SNP mutations on gene expressions, we treat the genotype data as covariates in this example. We care more about the information contained in the genotype data rather than the specific mutation status of each SNP. It is thus plausible to use the top principal components of the genotype data as a surrogate of the original ultra-high dimensional data in the SIFA model. However, if one is interested in integrating the genotype data and the gene expression data, raw data should be used. In that scenario, both data sets are primary data sets. we remark that one caveat of using the raw genotype data in that case is that the data may only take binary values or ternary values. It does not satisfy the Gaussian

assumption of the SIFA model. Extending the model to accommodate non-Gaussian data calls for more investigation.

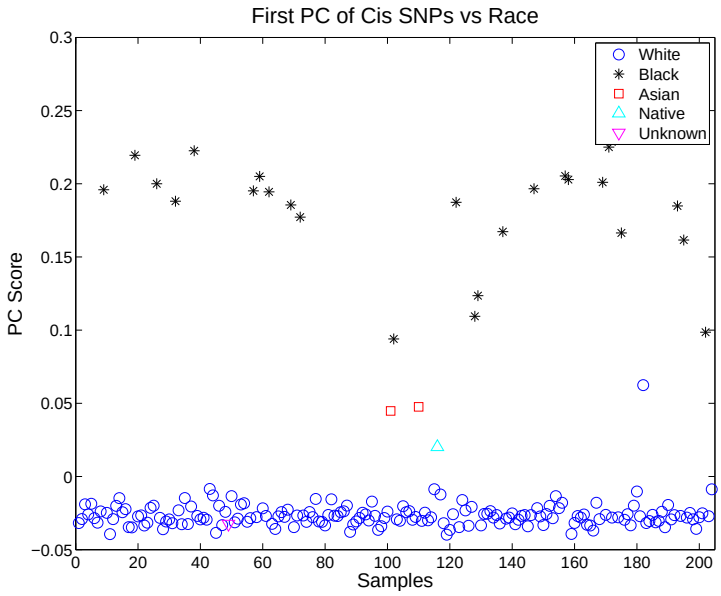

Figure S17: GTEx Example: the first PC score of the cis SNP data and the race information. The majority of the population is white (about 85%); the second largest group is African American (about 14%); some individuals are Asian or native American, or have unknown information.

E.2 Rank Estimation

In order to apply SIFA, we need to estimate the ranks for the joint structure and individual structures first. By visually inspecting the scree plot of the singular values of each data set, we do not observe any clear elbow point. Therefore, we exploit a variance explained criterion to estimate the signal ranks of different data sets in the first step of the two-step procedure described in the main manuscript. By setting the threshold to be 90%, we estimate of the intrinsic ranks for different data matrices to be $r^\star_{total} = 76, r^\star_1 = 50, r^\star_2 = 31, r^\star_3 = 46$. The actual proportions of variance explained against the ranks in different data sets are shown in Figure S18. Then by solving the equation system, in the second step, we obtain the joint and individual ranks for SIFA as $r_0 = 26$, $r_1 = 24$, $r_2 = 5$, and $r_3 = 20$. Since the ranks are all very large and each rank only explains a moderate amount of variation in each data set, it is reasonable

to perceive that the result is insensitive to minor changes of the ranks. Thus we omit the refining step through LCV. Note that the individual rank for blood ($r_2 = 5$) is much smaller than that for muscle or skin. To some extent, this is consistent with the previous GTEx finding that blood is an outlier tissue (Ardlie et al., 2015).

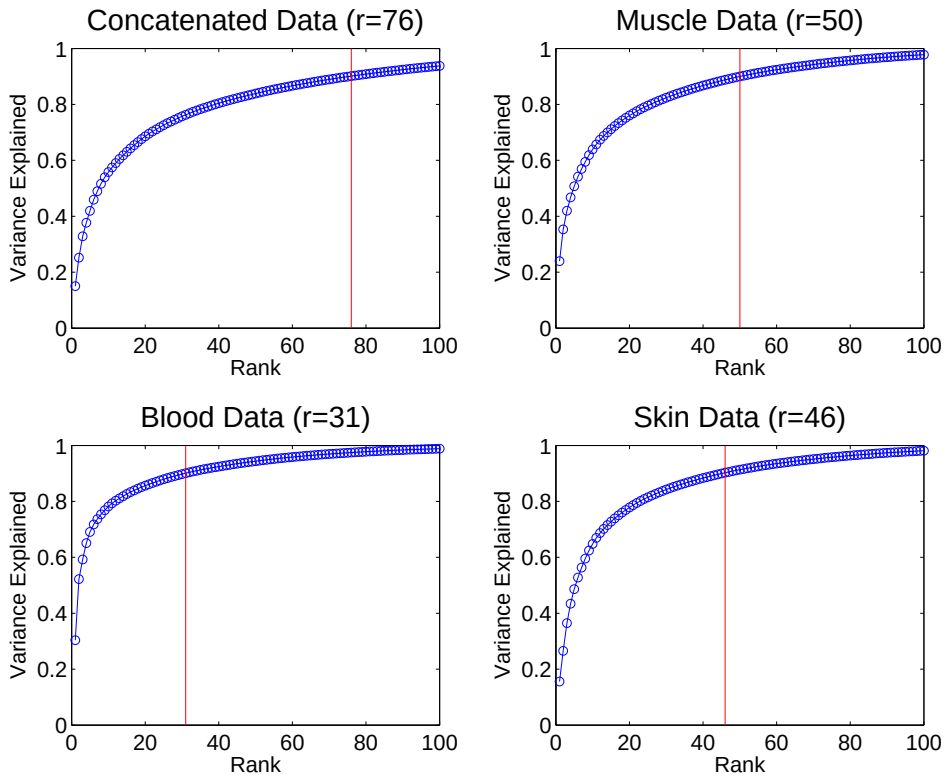

Figure S18: GTEx Example: The variance explained by top eigenvalues in different data sets (top left: concatenated data; top right: muscle data; bottom left: blood data; bottom right: skin data). The vertical red line in each plot indicates the smallest rank at which the variance explained exceeds 90% of the total variance.

E.3 Model Fitting

We fit a SIFA-B model to the data with linear relations between the covariates and the latent factors, for the following reasons. First of all, according to the preprocessing procedure, different primary data matrices have comparable scales (actually

identical in this data example). It is reasonable to assume that different data contribute equally to the joint factors, and thus the norm constraint in Condition B1 holds. Second, the high dimensionality of the gene expression data justifies the orthogonality constraints in Condition B1 and B2. Third, the linear relations between the covariates and the factors lead to high interpretability. The orthogonality between different components and the linear models make it very easy to decompose the total variation of the primary data into unrelated parts. Last but not least, from a computational perspective, the SIFA-B model with linear relations is extremely fast to fit (the algorithm converges with high precision within 1 minute on a standard desktop computer). All in all, we fit an SIFA-B model with linear relations to the data.

F Berkeley Growth Study

In this section, we apply SIFA to the Berkeley Growth Study data set (Ramsay and Silverman, 2002). The growth data consist of the height measurements of 39 boys and 54 girls from age 1 to 18. We particularly focus on the growth rate curves derived from the height measurements. As a preliminary analysis, we apply the conventional functional PCA (FPCA) (Ramsay and Silverman, 2002) to the data set. The first two principal loadings and scores are shown in Figure S19. The FPCA provides a low-rank representation of the data, but the quality of the dimension reduction is questionable. From the loading plots, we can see that the phase information and the amplitude information of the growth rates are highly confounded. The scatter plot of the scores shows there is a strong sex effect which is not incorporated into the FPCA. In addition, the two score vectors are highly associated with each other in a nonlinear way. All in all, the FPCA does not fully capture the underlying structure of the data, and may lead to misleading interpretation of the pediatric growth patterns.

To provide more effective dimension reduction and to better understand the population growth patterns, we decouple the original growth rate data into two related data sets, representing the "amplitude" and "phase" of growth respectively. This decomposition exploits a time-warping alignment of the observed functions to a common mean. The time-warping functions become the phase data set $\mathbf{Y}_2$, and the resulting aligned growth rate functions form the amplitude data set $\mathbf{Y}_1$. See Lee and Jung

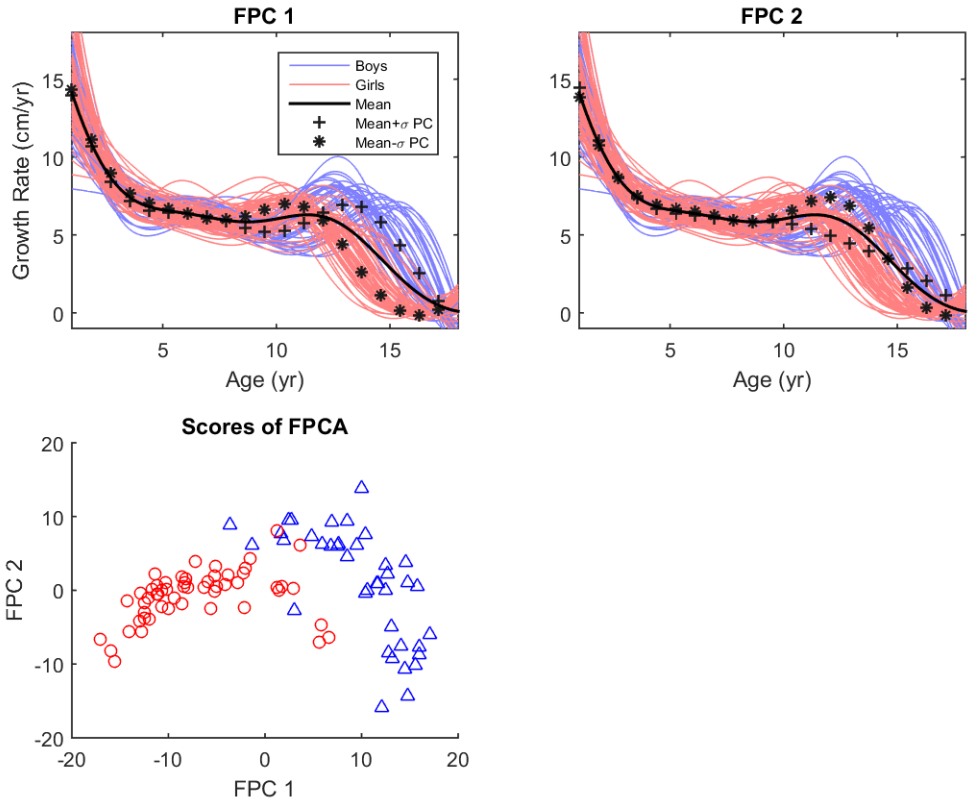

Figure S19: Berkeley Growth Rate Example: FPCA results. Top panels: the raw growth rate curves overlaid with the mean curve and the mean curve plus/minus one unit of the standard deviation of FPC (FPC1 on the left, and FPC2 on the right). Bottom panel: the scatter plot of FPC scores. Boys (blue triangles) and girls (red circles) are marked differently.

(2016) for a detailed exposition on the decomposition. Then we apply SIFA to the dual data sets $(\mathbf{Y}_1, \mathbf{Y}_2)$ jointly, and treat sex X as the univariate covariate. In particular, we fit the model under the general conditions, and exploit linear functions to capture the relations between sex and the latent factors.

By fixing the threshold for variance explained to be 90% and following the two-step procedure for rank estimation discussed in the main article, we obtain $(r_0, r_1, r_2) = (2, 2, 4)$. We model each data set as

$$\mathbf{Y}_k = \sum_{j=1}^{r_0} (\beta_{0,j} X + F_{0,j}) V_{0,j}^T + \sum_{j=1}^{r_k} (\beta_{k,j} X + F_{k,j}) V_{k,j}^T + \boldsymbol{E}_k, \ k = 1, 2.$$

When estimating the coefficient parameter $\beta_{k,j}$, we use a LASSO estimator (Tibshirani, 1996) to incorporate variable selection. Table S3 lists the estimated coefficients $\beta_{k,j}$, and the standard deviations for the random signals in $\boldsymbol{F}_0, \boldsymbol{F}_1, \boldsymbol{F}_2$ and the random noise in $\boldsymbol{E}_1, \boldsymbol{E}_2$. For better interpretation, the estimated loadings are shown in the original functional space in Figure S20, in a way similar to depicting the FPCA loadings. In addition, the conditional expectations of the latent factors are plotted against each other to visually understand the quality of the dimension reduction.

In this example, sex is the only covariate. The mean effect of sex is non-trivial for a joint component and one for each individual component, as shown in Table S3. To highlight the sex effect on pediatric growth patterns, the components explaining most of the sex effect are shown in Figure S20. In particular, we present the second joint component, the first amplitude component, and the third phase component. It is well known that pubertal growth spurt appears differently in boys and girls (both the intensity and the timing). The amplitude component plot and the phase component plot in the upper middle and right panels of Figure S20 confirm this. Girls (with negative covariates and negative coefficients for both components) tend to have a lower growth rate peak at a younger puberty age than boys. The joint component plot in the upper left panel shows an interesting joint pattern of amplitude and phase: girls (with negative covariates and a positive coefficient) tend to have a smaller and shorter gap between pre-pubertal dip and the pubertal peak compared to boys. The conventional FPCA fails to capture these patterns.

Some scatter plots for the conditional expectations of the SIFA factors are shown in the bottom panels of Figure S20. The scores are roughly elliptically distributed,

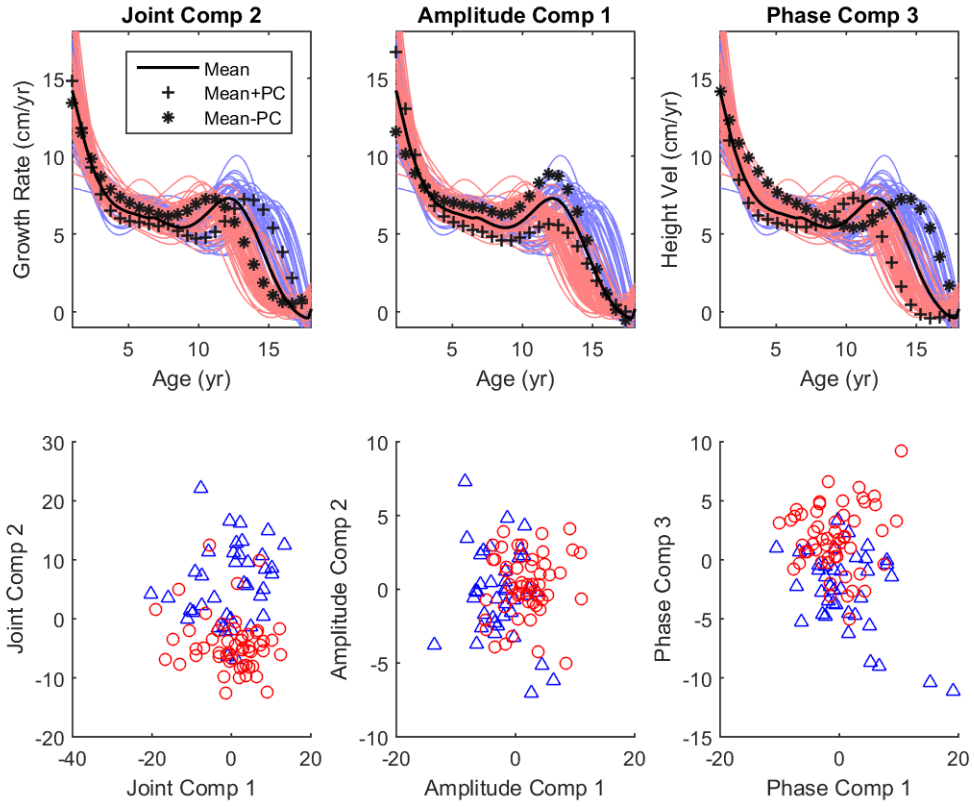

Figure S20: Berkeley Growth Rate Example: SIFA results. Top panels: the raw growth rate curves overlaid with the mean curve and the mean curve plus/minus one unit of the second joint loading (left), the first amplitude loading (middle), and third phase loading (right). Bottom panels: scatter plots of the first and second joint factors (left), the first and second amplitude factors (middle), and the first and third phase factors (right). Boys (blue triangles) and girls (red circles) are marked differently. Here the joint and individual components are selected based on how much they capture the sex effect.

Table S3: Berkeley Growth Rate Example: Estimated parameters of the SIFA model under the general conditions. The estimated coefficients for the sex effect are denoted by $\widehat{\beta}$ (Boys have positive covariate values and girls have negative covariate values). The estimated standard deviations of the random signals are denoted as $\widehat{\sigma_F}$. The estimated standard deviations of the random noise in different data sets are denoted as $\widehat{\sigma_E}$.

Component	$\widehat{\beta}$	$\widehat{\sigma_F}$	$\widehat{\sigma_E}$
Joint 1	0.00	7.30	
Joint 2	10.01	5.57	
Amplitude 1	-4.31	3.95	0.15
Amplitude 2	0	2.51	0.15
Phase 1	0	5.19	0.41
Phase 2	0	4.7	0.41
Phase 3	-4.37	3.04	0.41
Phase 4	0	2.13	0.41

indicating that there is no obvious association between different components. Therefore, when used in further statistical analyses such as prediction and inference, the SIFA factors may yield better results than the FPCA scores.